

Dealings on the Dark Web: An Examination of the Trust, Consumer Satisfaction, and the Efficacy of Interventions Against a Dark Web Cryptomarket

Vincent Harinam

Trinity Hall

August 2021

This thesis is submitted for the degree of Doctor of Philosophy.

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the Law Degree Committee.

Dealings on the Dark Web: An Examination of the Trust, Consumer Satisfaction, and the Efficacy of Interventions Against a Dark Web Cryptomarket

Vincent Harinam

Abstract

Objective. The overarching goal of this thesis is to better understand not only the network dynamics which undergird the function and operation of cryptomarkets but the nature of consumer satisfaction and trust on these platforms. More specifically, I endeavour to push the cryptomarket literature beyond its current theoretical and methodological limits by documenting the network structure of a cryptomarket, the factors which predicts for vendor trust, the efficacy of targeted strategies on the transactional network of a cryptomarket, and the dynamics which facilitate consumer satisfaction despite information asymmetry. Moreover, we also aim to test the generalizability of findings made in prior cryptomarket studies (Duxbury and Haynie, 2017; 2020; Norbutas, 2018).

Methods. I realize the aims of this research by using a buyer-seller dataset from the Abraxas cryptomarket (Branwen et al., 2015). Given the differences between the topics and the research questions featured, this thesis employs a variety of methodological techniques. Chapter two uses a combination of descriptive network analysis, community detection analysis, statistical modelling, and trajectory modelling. Chapter three utilizes three text analytic strategies: descriptive text analysis, sentiment analysis, and textual feature extraction. Finally, chapter four employs sequential node deletion pursuant to six law enforcement strategies: lead k (degree centrality), eccentricity, unique items bought/sold, cumulative reputation score, total purchase price, and random targeting.

Results. Social network analysis of the Abraxas cryptomarket revealed a large and diffuse network where the majority of buyers purchased from a small cohort of vendors. This theme of preferential selection of vendors on the part of buyers is repeated in other findings within this study. More generally, the Abraxas transactional network can then be viewed as set of transactional islands as opposed to a large, densely connected conglomeration of vendors and buyers. With regard buyer feedback, buyers are generally pleased with their transactions on Abraxas as long as the product arrives on time and is as advertised. In general, vendors have a relatively low bar to achieve when it comes to satisfying their customers. Based on the results of the sequential node deletion, random targeting was found to be ineffective across the five outcome measures, producing minimal and a slow disruptive effect. Finally, these strategies are based on a power law where a small percentage of deleted nodes is responsible for an outsized proportion of the disruptive impact.

Conclusion. As with all applied research examining emergent phenomena, this thesis lends itself to a more refined understanding of dark web cryptomarkets. While the results and conclusions drawn from these results are not perfectly generalizable to all cryptomarkets, they should serve to inform law enforcement on the dynamics which undergird these markets. To this extent, a sombre consideration of trust, consumer satisfaction, and tactical effectiveness of interventions is a necessary step towards the development of more effective countermeasures against these illicit online marketplaces. For law enforcement to be more effective against cryptomarkets, it is advised that an evidence-based approach be taken.

Table of Contents

Introduction	9
Dissertation Structure and Chapter Overview.....	10
Data and Methodological Overview.....	12
Conclusion	13
Chapter 1: Cryptomarkets: History, Operation, and Interventions	14
What is Cybercrime and How Organized is it?.....	14
What are Cryptomarkets?	17
Organizational Structure and Governance within Cryptomarkets	21
The Who, What, and Where of Cryptomarket Studies.....	28
Trust and Reputation	32
Law Enforcement Intervention and Network Disruption	39
Conclusion	43
Chapter 2: Trust under Uncertainty: How Network Structure and Vendor Selection Inform Trust Formation on Cryptomarkets	46
Literature Review	47
Trust in the Criminal World	47
Trust and Reputations on Cryptomarkets.....	48
Research Questions	50
Data.....	51
Methods	52
Descriptive Network Analysis.....	53
Community Detection Analysis.....	54
Variables and Model Estimations	54
Trajectory Modelling	56
Findings	57
Descriptive Statistics.....	57
The Network Structure of Abraxas	60
Community Detection Analysis.....	63
Regression Results and Power Few Distributions	67
Trajectory Results	69
Discussion	71

Trust, Reputation, and Network Structure on Abraxas	71
How is Trust Allocated on Abraxas?	72
Conclusion	73
Chapter 3: Consumer Satisfaction and Information Asymmetry on a Dark Web Cryptomarket: A Text Mining Approach	75
Literature Review	75
Customer Satisfaction in Online Reviews.....	75
Cryptomarkets and Consumer Satisfaction	77
Research Questions	78
Data.....	79
Methods	80
The Utility of Texting Mining	80
Variable Operationalization and Data Pre-processing	80
Descriptive Text Analysis	82
Sentiment Analysis	82
Feature Extraction: Logistic Lasso Regression	83
Findings	84
Descriptive Text Analysis	87
Sentiment Analysis (Ratings).....	91
Sentiment Analysis (Finalize Early)	92
Feature Extraction.....	93
Discussion	95
Finalize Early and Consumer Satisfaction	95
Information Asymmetry on Cryptomarkets	96
Conclusion	98
Chapter 4: Testing the Efficacy of Six Simulated Targeting Strategies on a Dark Web Cryptomarket	99
Literature Review	99
Typological Features of Criminal Networks	100
Simulated Interventions on Criminal Networks	101

Research Questions	103
Data	104
Methods	104
Simulation of Interventions Against Criminal Networks.....	105
Interventions that Target Only Network Capital	106
Interventions that Target Only Human Capital	106
Accounting for Network Adaptation.....	107
Findings	108
Simulation Results	108
Node Deletion Impact	112
Outcome Measure Carry-Over and Node Characteristics	113
Discussion	117
Targeting Based on Human and Network Capital	118
Metagames and Power Laws	119
Conclusion	120
Conclusion	121
Summary of Findings	122
Chapter 2	122
Chapter 3	123
Chapter 4	124
Implications for Dark Web Interventions	125
Bastions of Responsible Use	129
Theoretical Contribution to Criminology	130
Future Research.....	132
Appendix 1. Supplementary material for chapter 4	134
References	137
Acknowledgements	151

List of Figures

Chapter 2 Figures

Figure 1: Abraxas Transactional Network	60
---	----

Figure 2: Abraxas Transactional Network by Community	64
Figure 3: Power Law Distributions of Vendors by Transactions, Buyers, and Revenue	68
Figure 4(a): K-Means Trajectories, Number of Transactions (Success)	70
Figure 4(b): K-Means Trajectories, Number of Unique Buyers (Popularity).....	70
Figure 4(c): K-Means Trajectories, Cumulative Revenue (Affluence)	70

Chapter 3 Figures

Figure 1(a): Word frequencies top 20 (all)	87
Figure 1(b): Word frequencies top 20 (> 5's)	87
Figure 1(c): Word frequencies top 20 (< 5's)	87
Figure 2: Power law distribution for word frequency	88
Figure 3(a): Cluster dendrogram (all).....	89
Figure 3(b): Cluster dendrogram (<5's)	89
Figure 3(c): Cluster dendrogram (5's).....	89
Figure 4(a): Sentiment polarity (all).....	91
Figure 4(b): Sentiment polarity (<5's).....	91
Figure 4(c): Sentiment polarity (5's).....	91
Figure 5: Sentiment cloud; 5's and <5's.....	92
Figure 6(a): Sentiment polarity (FE)	93
Figure 6(b): Sentiment polarity (No FE)	93
Figure 7: Error as a function of lambda.....	94

Chapter 4 Figures

Figure 1(a): Number of Components.....	110
Figure 1(b): Number of Isolates	110
Figure 1(c): Number of Nodes in Components	111
Figure 1(d): Number of Nodes in Largest Components	111
Figure 1(e): Average Geodesic Distance	111
Figure 2(a): Node Removal Impact - Number of Nodes in Largest Components ...	112
Figure 2(b): Node Removal Impact - Number of Isolates	112
Figure 2(c): Node Removal Impact - Number of Components	113

Figure 2(d): Node Removal Impact – Average Nodes in Components.....	113
Figure 2(e): Node Removal Impact - Average Geodesic Distance	113

List of Tables

Chapter 2 Tables

Table 1: Descriptive statistics of variables used in analysis	56
Table 2: Descriptive Statistics on the Abraxas Cryptomarket	57
Table 3: Network characteristics	60
Table 4: Distribution of Network Components	61
Table 5: Frequency of Unique Vendors Purchased from by Number of Transactions	62
Table 6: Distribution of In and Out Degree	62
Table 7: Community Network characteristics	64
Table 8: Community Network Measures (Top 20 based on community size)	65
Table 9: Communities by Item Categories and Country Shipped From (Top 20 based on community size)	66
Table 10: Results of Regression Models	67
Table 11: Summary of K-Means Trajectories.....	69

Chapter 3 Tables

Table 1: Descriptive statistics (ratings by product type and finalize early).....	85
Table 2: Word and character statistics	86
Table 3: Word frequency top 5 by product type and rating	89
Table 4(a): Word Associations (5's)	90
Table 4(b): Word Associations (<5's)	90
Table 5: Distribution of sentiment polarity (ratings).....	91
Table 6: Distribution of sentiment polarity (finalize early)	93
Table 7: Logistic lasso coefficients and odds ratios	94

Chapter 4 Tables

Table 1: Network characteristics	104
Table 2: Impact of Single Node Deletions by Strategy and Outcome	108
Table 3: Top 100 Actors Held in Common Across Targeting Strategies	114

Table 4: Descriptive Statistics for Top 100 Actors Held in Common Across Targeting Strategies..... **115**

Introduction: What's to Come

Gone are the days when prospective consumers need rely solely on local dealers to procure drugs and other illicit goods and services. The advent of digital encryption and internet connectivity has facilitated the rise of cryptomarkets. Similar to Amazon or eBay, these are illicit online marketplaces hosted on the dark web which facilitate the truck, barter, and trade of illegal goods and services. Much like licit online markets, cryptomarkets permit those seeking to purchase illicit goods and services to do so from the comfort of their own home, placing their order with a vendor and receiving the product through the postal service. Be it marijuana, cocaine, bladed implements, or hitmen, these platforms are replete with a variety of illicit wares.

Cryptomarkets represent a unique permutation that both improves upon traditional criminal dynamics while introducing new elements that challenge the capabilities of law enforcement. Moreover, these platforms present a novel opportunity for researchers to test the accuracy of key theoretical precepts that are present in terrestrial markets. How are trust and reputation associated with the network structure of a cryptomarket? How is information asymmetry mitigated or overcome and what can we learn from it? What factors create and sustain consumer satisfaction? What are the structural vulnerabilities in cryptomarket transactional networks? Which strategic interventions initiated by law enforcement work best? How do these strategic interventions differ in their stated objective and measured outcomes? These are some of the questions which will be investigated in the forthcoming chapters.

To this extent, the overarching goal of this dissertation is to better understand not only the network dynamics which undergird the function and operation of cryptomarkets but the nature of consumer satisfaction and trust on these platforms. More specifically, I endeavour to push the cryptomarket literature beyond its current theoretical and methodological limits by documenting the network structure of a cryptomarket, the factors which predicts for vendor trust, the efficacy of targeted strategies on the transactional network of a cryptomarket, and the dynamics which facilitate consumer satisfaction despite information asymmetry. Moreover, I also aim to test the generalizability of findings made in prior cryptomarket studies (Duxbury and Haynie, 2017; 2020; Norbutas, 2018). This thesis utilizes several methodological techniques to answer the various research questions it posits, leveraging a combination of social network analysis, statistical modelling, text mining, and adaptive computer simulations.

The specific aim of this dissertation is two-fold. First, I seek to push the theoretical boundaries of cryptomarket research in order to better understand the functional mechanisms of cryptomarkets. That is, I will use cryptomarkets as a testbed for social scientific theories that propose conditions under which anonymous actors are more likely to trust each other, and the mechanisms that increase cooperation under uncertainty. While the technology that allows cryptomarkets to operate in the manner that they do is certainly important, I am primarily interested in the network dynamics between participants and how these affect the overarching structure and robustness of these markets. Furthermore, computer-mediated interactions on the Internet provide new opportunities to examine the links between reputation, information asymmetry, and the development of trust between individuals who engage in various types of illicit exchange. While some researchers have dealt with some of topics featured in this thesis, cryptomarket scholars are uncertain about the generalizability of these findings given the novelty of this criminological phenomenon. Indeed, more research is

required in specific areas to better understand the function and operation of these illicit online marketplaces.

The second aim of this dissertation is to use the findings herein to inform targeted interventions by law enforcement against cryptomarkets. Past law enforcement strategies targeting cryptomarkets have been ineffective and, in some cases, counterproductive (Soska and Christin, 2015; Decary-Hetu and Giommoni, 2017; van Buskirk et al., 2017). As such, this thesis' explicit focus on trust dynamics, consumer satisfaction, and efficacy of law enforcement interventions might offer some insight into how law enforcement might structure their cryptomarket intervention strategies to achieve maximum long-term disruptive impact. By posing new questions and revisiting old ones, I seek to explain how cryptomarket participants engage with one another despite the limitations of information asymmetry and how this affects consumer satisfaction and the structure of a cryptomarket's transactional network.

Dissertation Structure and Chapter Overview

This dissertation is structured around four disconnected chapters, with the first serving as an up-to-date consolidation of the cryptomarket literature and the second, third, and fourth chapters addressing a distinct set of research questions pertaining to a specific topic that is unaddressed or partially examined within the extant literature. To this extent, chapters two, three, and four will focus, in order, on: 1) the network structure and trust dynamics of a cryptomarket, 2) the elements which predict for consumer satisfaction or dissatisfaction on a cryptomarket, and 3) the efficacy of six targeting strategies in disrupting a cryptomarket's ease of operation. Each of these topics and their associated research questions were selected after an extensive examination of the cryptomarket literature. Indeed, they both represent a critical gap in the scholarly literature and function as a key pedagogical hurdle that must be overcome for cryptomarket research to progress further.

Chapter 1 is an up-to-date summary of the extant cryptomarket literature, drawing upon a vast swath of studies across a decade of research. As such, there are no research questions posed or analyses conducted in this chapter. The objective of this chapter is to both explain what cryptomarkets are and situate these illicit platforms within the cybercrime and organized crime contexts. A secondary objective is to take stock of the current state of cryptomarket research, tracking major scholarly themes across a decade of research. To this extent, this chapter will be separated into six sections: 1) what is cybercrime and how organized is it?, 2) what are cryptomarkets?, 3) the organizational structure and governance within cryptomarkets, 4) the who, the what, and the where of cryptomarket studies, 5) trust and reputation on cryptomarkets, and 6) law enforcement interventions and network disruptions of cryptomarkets.

The body of this thesis will consist of three distinct (though interrelated) research papers that cover a separate area of inquiry. Following Duxbury and Haynie (2017) and Norbutas (2018), chapter two examines the network structure of a cryptomarket. More specifically, I seek to identify the market-level metrics that predict for vendor selection as well as the developmental trajectory of vendor trustworthiness. In short, this chapter seeks to disentangle the overarching concept of trust on cryptomarkets by both revisiting the findings made in prior studies (Duxbury and Haynie, 2017; Norbutas) and generating new findings using new conceptualizations and methods. This chapter seeks to replicate findings relating to the network structure of cryptomarkets made in prior studies. It will, however, contribute new

material to the literature by examining new predictors across three conceptual definitions of vendor trustworthiness. This will also include an examination of the developmental trajectory of vendor trustworthiness; a first within the cryptomarket literature. Importantly, the Abraxas cryptomarket will be examined. Chapter 2 answers four research questions:

1. What is the network structure of Abraxas?
2. What is the composition of transactional communities within the network?
3. What market-level metrics and/or vendor characteristics predict for vendor trustworthiness (i.e. success (completed transactions), popularity (unique buyers), and affluence (revenue))?
4. What is the developmental trajectory of vendors' success, popularity, and affluence during their tenure on Abraxas?

Chapter three seeks to identify and compare the determinants of customer satisfaction and dissatisfaction among buyers on a cryptomarket. This is the first such study to both examine the lexical predictors of vendor ratings as well as the sentiment structure of qualitative reviews. As such, there is an explicit focus on determining the similarities and differences between five-star and non-five-star ratings and how this might affect information asymmetry on dark web markets. Additionally, I examine role of “finalizing early” in mitigating information asymmetry. While previous studies (Hardy and Norgaard, 2016; Janetos and Tilly, 2017; Przepiorka, Norbutas, and Corten, 2017; Tzanetakis, Kamphausen, Werse, and von Laufenberg, 2016) have examined the impact of dark market rating systems on vendor success and profitability, none have examined this phenomenon using textual data. Moreover, there has been no research on the factors that affect the written reviews buyers leave for vendors and whether and how consumers' attitudes affect their overall ratings of vendors. Chapter 3 will continue to focus on Abraxas, answering three research questions:

1. Based on written reviews, what are the determinants of consumer satisfaction and dissatisfaction among buyers on Abraxas?
2. Does the sentiment structure of positive and negative reviews differ? If so, to what extent?
3. What words best predict five and non-five ratings among buyers?

Finally, chapter four examines the efficacy of six law enforcement targeting strategies: lead k, eccentricity, total revenue generated, cumulative reputation score, listing amount, and random targeting. To this extent, sequential node deletion will be utilized. Five outcome variables (number of isolates, number of components, average number of nodes in components, average geodesic distance, and number of nodes in the largest component) are used to measure the efficacy of each targeting strategy. The study seeks to test the generalizability of Duxbury and Haynie's (2018; 2020) findings on a different cryptomarket, Abraxas. More importantly, however, this study is the first to answer questions regarding the similarity of targeting strategies as well as their short and long-term efficacy. It will serve as the most in-depth examination of strategic interventions against cryptomarkets. Whereas several studies (Xu and Chen, 2003; Keegan et al., 2010) have failed to incorporate network adaption and preferential selection processes into their simulations, this study will set parameters to govern the (purported) behaviour of actors when nodes are removed. This chapter answers three research questions:

1. Of the six proposed disruption strategies, which offers the greatest initial amount of damage to a criminal network?
2. Of the first 100 nodes that are removed per each disruption strategy, does their impact carry-over across all outcome measures?
3. What do these strategies tell us about the efficacy of dark web disruption strategies?

Importantly, as chapters two, three, and four are individual research papers containing their own distinct literature reviews, there will be some overlap between portions of the first chapter and portions of the literature reviews in each succeeding paper. Nevertheless, each chapter offers unique insight into the functional mechanisms which govern transactional exchanges on cryptomarkets between buyers and vendors.

Data and Methodological Overview

I realize the aims of this research by using a buyer-seller dataset from the Abraxas cryptomarket (Branwen et al., 2015). Apart from the anonymous cryptomarket analysed by Duxbury and Haynie (2017), this is the only marketplace where unique identifiers are available for buyers. As such, it was the only known publicly available dataset which allowed for network analysis and adaptive computer simulation. With assistance from Lukas Norbutas of Utrecht University and Cambridge University's Computer Laboratory, this data was extracted from a public data repository established by independent researcher, Gwern Branwen. This data repository contains scraped webpages from 2013 to 2015. Given the infrequent nature of the scrapes, not all webpages have been collected. Nevertheless, Norbutas (2018) estimates that crawls of Abraxas have successfully collected 92.4% of all listed items on the Abraxas cryptomarket. This includes information on vendor name, vendor shipping location, listing title, listing price, listing description, transaction date, buyer unique identifier, buyer rating, and buyer feedback.

HTML links in the dataset were restitched together in Python to recreate the Abraxas website. Thus, this recreated website serves as a copy of the original Abraxas cryptomarket, possessing information on transactions that were successfully scraped. Furthermore, each webpage in the dataset was manually inspected to identify duplicate transactions based on the feedback provided. While buyers might leave feedback on their original post, they may return to alter the message. As such, extracting data from these webpages would yield duplicate transactions if each transaction was not properly inspected. Once all duplicates were identified and removed, I was left with a total of 5434 transactions over a period of 7 months (January to July) in 2015. These were stored in an Excel spreadsheet. While Abraxas was established in December of 2014, the first transaction occurred on January 15th of 2015. It is important to clarify that this dataset does not include all recorded transactions on Abraxas. This is due to both the infrequency of the scrapes conducted by Branwen (2015) and the vast number of broken webpages that could not be repaired and accessed. As such, while this dataset includes numbers sufficient for analysis, it does not include the full cohort of transactions on the cryptomarket. This is a clear limitation.

Nevertheless, there were 269 unique sellers and 2794 unique buyers in the dataset. Importantly, the Abraxas dataset was previously used by Norbutas (2018) in an examination of the geographical distribution of transactions. For my purposes, I reconstruct a two-mode buyer-seller trade network. These data were used in chapters two and four while chapter three utilized written feedback provided by buyers from each successful transaction.

Given the differences between the topics and the research questions featured in chapters two, three, and four, this thesis employs a variety of methodological techniques. Chapter two uses a combination of descriptive network analysis, community detection analysis, statistical modelling, and trajectory modelling. Descriptive statistics were used to summarize market transactions. This is done to understand both the nature and composition of illicit transactions on Abraxas. In contrast, community detection analysis is used to discern the subgroup structure of this transactional criminal network. As well, three regression models were used to determine the predictors of vendor trustworthiness. To measure vendor trustworthiness, three proxy variables were created: success, popularity, and affluence. As trust is manifested in a variety of ways, each of these dependent variables reflects a key element of trust. Finally, this chapter leverages k-means trajectory modelling to examine the developmental pattern of vendor trustworthiness on Abraxas.

Chapter three utilizes three text analytic strategies: descriptive text analysis, sentiment analysis, and textual feature extraction. All analyses and visualizations were conducted in R. Descriptive text analysis is a fairly standard text mining procedure. Simple term frequencies are conducted to identify the words used by Abraxas buyers to describe their experience. Furthermore, sentiment scoring is conducted on the written reviews. Sentiment scoring measures the positive or negative intent in a writer's tone. Finally, feature extraction is used to understand what words predict for customer satisfaction and dissatisfaction. To this effect, a supervised machine learning technique, logistic lasso regression, is utilized.

Chapter four employs sequential node deletion pursuant to six law enforcement strategies: lead k (degree centrality), eccentricity, unique items bought/sold, cumulative reputation score, total purchase price, and random targeting. Five outcome variables (number of isolates, number of components, average number of nodes in components, average geodesic distance, and number of nodes in the largest component) are used to measure the impact of each targeting strategy. This study sets parameters to govern the purported behaviour of actors when nodes are removed. As such, the transactional network's overall behaviour can be accurately modelled (Bright et. al, 2018) through an evidence-based calculus.

Conclusion

As with all applied research examining emergent phenomena, this thesis lends itself to a more refined understanding of dark web cryptomarkets. More importantly, the following chapters were conceptualized, developed, and written with the sole intent of improving current law enforcement strategies which target cryptomarkets. While the results and conclusions drawn from these results are not perfectly generalizable to all cryptomarkets, they should serve to inform law enforcement on the dynamics which undergird these markets. To this extent, a sombre consideration of trust, consumer satisfaction, and tactical effectiveness of interventions is a necessary step towards the development of more effective countermeasures against these illicit online marketplaces. For law enforcement to be more effective against cryptomarkets, it is advised that an evidence-based approach be taken.

Chapter 1: Cryptomarkets: History, Operation, and Law Enforcement Interventions

This chapter serves to consolidate the scholarly literature on cryptomarkets, identifying and explicating all strands of scholarly work on this topic. As such, this chapter will function as an extended literature review, distilling findings from peer-reviewed works while offering a measured examination of cryptomarkets within the context of the cybercrime and organized crime literatures. Moreover, this chapter will consist of six sections. First, I detail the phenomenon of cybercrime, exploring its origin, transformation, and the extent to which it is organized. Second, I explore the genesis and general operation of cryptomarkets, examining the importance of onion routing and cryptocurrencies. I then examine the organizational features of cryptomarkets. Here, I highlight the hierarchical administrative structure of cryptomarkets as well as the mode of governance and flexible exchange networks imbedded within. Fourth, I explore the three primary strands of cryptomarket research which detail the participants, countries, and products featured on these platforms. Fifth, I consider the role of trust and reputation on cryptomarkets, detailing the various mechanisms used by vendors to instil trust in buyers. Finally, I close out this chapter by exploring the various actions taken by law enforcement organizations against cryptomarkets. This will also include research on simulated interventions against the transactional network of these illicit entities.

What is Cybercrime and How Organized is it?

The volume and sophistication of cybercrime operations have dramatically increased in the last decade (Holt, Bossler, and Malinski, 2016). Fraudsters are exploiting email systems to ensnare unassuming victims with faulty services and get-rich quick schemes (Grabosky, 2007), internet chatrooms and message boards are being used to solicit sex and, in some cases, prop up the international sex trade (Farley, Franzblau, and Kennedy, 2013), and social media platforms are being used by youth to bully their classmates (Hinduja and Patchin, 2012). In addition, technology has given birth to entirely new forms of crime. Distributed denial of service attacks and malicious software are two such computer-assisted offenses that have produced substantial economic harm (Bossler and Holt, 2012; Holt and Turner, 2012). Each day brings new challenges for law enforcement within the realm of cyber.

Though there exists no formal definition, scholars generally agree that cybercrime “involves the use of cyberspace or computer technology to facilitate acts of crime and deviance” (Bossler and Holt, 2016, 45). Moreover, Grabosky (2007) categorizes cybercrime along three conceptual dimensions: computers as the instrument of crime, computers as the target of criminal activity, and computers as incidental to criminal activity. Nevertheless, this categorization falls in short in one respect. While this classification system creates conceptual boundaries, it is often subject to categorical overlap. In short, certain cyber-enabled crimes can fall within multiple categories. Consider botnets. These are networks of infected computers that are remotely controlled by another computer (Ianelli and Hackworth, 2005). In this case, computers are both the instrument and the target of the offence.

However, Wall (2001) subdivides cybercrime into four categories: cyber-trespass, cyber-deception, cyber-pornography, and cyber-violence. Much like trespass in an offline setting, cyber-trespass involves accessing a computer system without the expressed consent of the owner. Similarly, cyber-deception, the second category, involves the use of the internet

to illegally acquire information from people or corporations. Importantly, cyber-trespass and cyber-deception are fundamentally linked to the concept of hacking. Holt and Bossler characterize hackers as “individuals who create viruses and botnet codes which lead to automated malicious attacks and/or actively participate in attacks against computer systems and sensitive networks” (2014, 22). Indeed, the authors’ primary contention is that hackers are best conceptualized as criminals and deviants. Brewer and Goldsmith, in contrast, attempt to establish the moral and legal versatility of hackers in proposing the term “digital drift”. The authors maintain that “new technologies enable individuals to both ‘embed’ and ‘dis-embed’ themselves in a variety of criminal activities and lifestyles off- as well as online” (2015, 113). As such, hackers do not perpetually lead a life of crime but instead drift between periodic stints of cyber-criminality and obedience to the law.

The third category, cyber-pornography “encompasses the range of sexual expression enabled by computer-mediated communications and the distribution of sexually explicit materials on-line” (Wall, 2007, 32). This particular category is the most controversial of the four. Indeed, online pornography is not an illegal activity in and of itself but is rather a feature of the internet, representing a large proportion of internet traffic. Nevertheless, cyber-enabled child pornography and sexual exploitation are crimes which might better fit this category. The fourth category in Wall's (2001) cybercrime typology is cyber-violence. This refers to actions taken by individuals which harm others in both online and offline settings. This generally includes stalking, harassment, and bullying online. There is, nevertheless, a glaring conceptual problem with this category. In particular, the use of term “violence” is not wholly descriptive of crimes within this category. More specifically, whereas violence typically constitutes a physical action causing bodily harm, online stalking, harassment, and bullying are not themselves physical acts. Each of these crimes occur in cyberspace and do not allow for physically harm against the victim. In general, Wall’s (2001) categorizations, while helpful at the time of its conception, are not satisfactory in a contemporary cybercrime setting. Indeed, the constantly changing nature of cybercrime renders definitions and categorizations obsolete over time.

Of critical importance then is the larger debate surrounding the novelty of cybercrime. Is it “old wine in new bottles” or “new wine in new bottles” (Grabosky and Smith, 2001; Wall, 1999; Wall, 2007; Yar, 2005)? That is to say, are cybercrimes merely terrestrial crimes that have taken a different form or are they an entirely new permutation that is actualized in a different manner? Indeed, the composition of a crime committed in cyberspace is, by Grabosky’s (2001) estimation, congruent to those committed in a physical setting. To elaborate, cybercrime, like any terrestrial crime, can be explained by the intersection of three requisite factors: a suitable target, motivated offender, and lack of a capable guardian. This is referred to as routine activities theory (Cohen and Felson, 1979). Per Cohen and Felson’s (1979) theory, the infiltration of a medical database or distribution of malware must possess these qualities for it to have taken place (Grabosky and Smith, 2001). This is the same for the theft of a car or the murder of a rival gangster.

However, the application of routine activities theory to cybercrime is refuted by Yar (2005). To this extent, Yar (2005) contends that “whereas people, objects and activities can be clearly located within relatively fixed and ordered spatio-temporal configurations in the ‘real world’, such orderings appear to destabilize in the virtual world” (424). As such, one cannot easily extrapolate the precepts of routine activities theory to cybercrime. Indeed, one

of the key characteristics which separates digital criminality from terrestrial criminality is the potential for transnational offending. Many, if not most, cybercrimes can now take place in one jurisdiction but be initiated in another. This presents serious complications if the laws and priorities in each involved jurisdiction differs. If a citizen of one country were to fall victim to an online investment fraud originating in another, both or neither of the authorities in each involved nation may have an investigative or punitive interest.

Regardless of its novelty or lack thereof, the organizational structure of cybercrime is a matter that has been rigorously discussed by scholars (Grabosky, 2007; Lusthaus, 2012; Wall, 2001; Wall, 2015). Since cybercrimes are the product of networked computers, they have fundamentally transformed the scale and efficiency of criminal operations. These transformations have given rise to sophisticated organizations that are locally-hosted yet globally-active. More importantly, the use of technology contributes to the reorganization of traditional divisions of labour within a criminal organization. On one hand, they serve to both automate and deskill certain criminal activities, while on the other reskilling and empowering individuals and groups to operate a criminal enterprise (Pease, 1991, 24; Savona and Mignone, 2004; Wall, 2007).

In this regard, Lusthaus (2018) makes the argument that cybercrime has evolved from mischievous activities carried out by disparate actors to a profit-driven industry that is dependent on anonymity. The historical evolution of hacking is demonstrative of this change. Between the 1950's and 1980's, hacking resided within the domain of scientific inquiry as university and government-backed researchers waded into the maliciousness of phreaking. This changed, however, in the 1990s with the proliferation of desktop computers. Individual hacking metathesized into organized trading forums which then gave way to professional groups that carried out coordinated attacks. As such, a growing level of collaboration paired with an increasing desire for professionalization created an economic infrastructure based around trust and anonymity.

Cybercrime operations, be they carding forums or hacker groups, function according to the same principles followed by industrial organizations (Lusthaus, 2018). That is to say, there are clear divisions of labour by which different activities, from hacking to coding, are handled by different specialists. Moreover, increasing specialization leads to increasing professionalization. By specializing in a specific activity or task within the cybercrime supply chain, participants are encouraged to both hone their creative talents and market them to willing customers and business partners. An influx of actors and associated firms creates more options for collaboration and networking. Under these conditions, monetization becomes inevitable. Profit and plunder have superseded past desires for fun and intellectual challenge (Grabosky, 2007; Lusthaus, 2018). This change in motivation puts additional strain on law enforcement as they must curtail the efforts of malicious, enterprising actors as opposed to those looking for a good time.

All told, the unique and enduring characteristic which typifies cybercrime, organized or otherwise, is its malleability. It is never a simple, stagnant operation or enterprise. Rather, it is a practice and activity that shifts in form and orientation depending on the technology available and the expertise required. In this regard, advances in digital cryptography and peer-to-peer monetary systems have allowed for the growth of illicit online marketplaces that have taken root on the dark web. These cryptomarkets, as they are called, are a unique

permutation that both improves upon traditional crime dynamics while introducing new elements that challenge the capabilities of law enforcement. I will detail the history and operation of these illicit platforms in the next section.

What are Cryptomarkets?

Though unaware of its distal impact, Peter Grabosky (2007) was one of the first to consider the looming possibility of emerging technologies amplifying the organized distribution of illicit goods and services. Of course, our propensity for using exponential technologies to engage in the truck, barter, and trade of illicit commodities is not new. The first official e-commerce transaction occurred in 1972 when students from MIT and Stanford University utilized ARPANET (a 1960's packet switching network that evolved into the Internet) to negotiate the sale and purchase of marijuana. From this inauspicious beginning, illicit online transactions have evolved at a rapid pace.

Buoyed by technological advancements, globalization, and market innovations, illicit goods and services are more readily accessible to those with the requisite know-how. In fact, recent research (Martin, 2014a) has indicated that illicit online markets have become “hybrid markets that combine traditional social and economic opportunity structures with newer opportunities provided by the internet. Not only has the internet created new avenues for criminal networking, but it has also reconfigured traditional relations among suppliers, intermediaries, and buyers” (56). These developments are punctuated within cryptomarkets. These entities are the culmination of many decades of innovation within the realm of cyber.

The operational history of cryptomarkets is rather brief, dating back to as early as 2011. The first cryptomarket, Silk Road, was founded in 2011 by the enigmatic Ross Ulbricht, a physics major from Austin, Texas (Martin, 2014a). However, this site was shut down in 2013 following the FBI's arrest of Ulbricht. In the succeeding months, new cryptomarkets began to emerge with Silk Road 2.0 (a direct successor), Agora, Atlantis, and CannabisRoad leading the way. These markets would be shut down in 2014 following Operation Onymous, a joint initiative by the NCA, FBI, and Europol (Decary-Hetu and Giommoni, 2017). By 2015, however, the markets would again readjust as AlphaBay became the most prosperous cryptomarket to date. Finally, in July 2017, Operation Bayonet, a joint operation by the by the FBI, DEA, Europol, and Dutch National Police, led to takedowns of AlphaBay and Hansa, the first and third largest cryptomarkets at that time. But the question remains: what are cryptomarkets?

James Martin defines a cryptomarket as “an online forum where goods and services are exchanged between parties who use digital encryption to conceal their identities” (2014a, 2). While Martin's (2014a) definition is the most popular, it is not without its flaws. In particular, this definition lacks a marked level of specificity, conflating illicit online marketplaces with forums. While online forums do cater to the trade of illicit goods and services by advertising them (Dupont, Cote, and Decary-Hetu, 2016; Hutchings and Holt, 2015), they are distinct podia designed primarily for discussion and debate. Functionally speaking, online marketplaces, licit or otherwise, do not permit for thread-based discussions. They are first and foremost marketplaces where goods and services are bought and sold.

This important feature is captured in a much-improved definition created by Barratt and Aldridge (2016). The authors define cryptomarkets as “marketplaces that host multiple sellers or ‘vendors’, provides participants with anonymity via its location on the hidden web and use of cryptocurrencies for payment, and aggregates and displays customer feedback ratings and comments” (Barratt and Aldridge, 2016, 78). Though lengthy, the strength of this

definition lies in its exactness. It outlines the various idiosyncrasies which help distinguish cryptomarkets from other illicit markets that exist in cyberspace as well as terrestrial settings. Cryptomarkets are thus characterized by their location on the dark web, use of cryptocurrencies and feedback systems, and hosting of buyers and vendors.

To this extent, cryptomarkets do not actually sell anything (Christin, 2013; Martin, 2014b). These illicit online marketplaces function more as brokerage platforms which bring together buyers and vendors willing to engage in voluntary economic transactions over a multitude of illicit goods and services. To this extent, Christin (2013) notes that cryptomarkets are risk management platforms for criminals. By eliminating physical interactions between transacting parties, cryptomarkets serve to reduce and, to an extent, eliminate the potential for physical violence (Barratt et al., 2016; Morselli et al., 2007). Moreover, the anonymity and escrow services embedded within a cryptomarket's transactional infrastructure reduces risk as it relates to fraudulent exchanges. Importantly, these methods aid in obfuscating the activities of cryptomarket participants, increasing the difficulty of law enforcement in identifying much less apprehending these actors.

Furthermore, the financial escrow system is particularly important as it mimics the financial risk reduction competencies of similar systems developed by licit electronic commerce platforms like eBay or Amazon. To this extent, Christin (2013) notes that an escrow service ensures that funds are kept until a transaction is "finalized" by the buyer and released to the vendor. Suppose Alice wanted to purchase an item from Paolo. Instead of paying Paolo directly, Alice would pay the marketplace operator, Manuel, who would then direct Paolo to ship the item to Alice. Once Alice confirms that she has received the item, Manuel would then release the money to Paolo while keeping a small fee for himself. This payment system allows cryptomarket operators to adjudicate any dispute that could arise should a vendor claim that an item had been shipped, but the buyer claims to have not received it. Nevertheless, a buyer may "finalize early" (FE), foregoing the escrow system and simply transferring the funds immediately upon purchase. The phenomenon of early finalization has not been examined extensively by cryptomarket scholars.

However, to truly understand what cryptomarkets are, it is important to situate these digital phenomena within their place of operation. The Internet, as we understand it, is segmented into two distinct parts: the surface web and the deep web. All content that is accessible via a search engine such as Google or Bing are part of the surface web. These websites are indexed by a search engine and are thus publicly accessible, requiring no special configurations or permission to access them. In contrast, web pages that are not indexed and accessible by a search engine are part of the deep web. According to Epstein (2014), the deep web is estimated to contain 96% of all networked webpages, making it nearly 500 times larger than the surface web. However, deep web content is for the most part legal. This includes "content that is locked behind paywalled websites, content accessible through company or academic databases, any kind of database that cannot be searched directly by Google, websites that are not linked to other websites, private websites and forums, and large amounts of social networking site content (e.g., non-public Facebook content)" (Barratt and Aldridge, 2016, 79).

Cryptomarkets, however, are situated in a small subset (a hidden overlay network) of the deep web called the dark web. In this regard, dark web internet services are inaccessible without unique configurations, explicit authorization, or a specialized browser (Barrett,

Lenton, Maddox, and Allen, 2015, 50; Gaup, 2008). As such, these websites are not indexed by a search engine and are not publicly accessible. However, this is not to suggest the inherent criminality or maliciousness of all actors operating on the dark web. Dark web platforms are often utilized by political dissidents and whistle-blowers seeking to bypass draconian censorship laws and government overreach (Bradbury, 2014; Hardy and Norgaard, 2016). Indeed, this was the original purpose of the software that permits access to the dark web.

Importantly, the feature which separates cryptomarkets from other illicit exchange networks and distribution systems is its reliance on encryption technology. As Decary-Hetu and Giommoni (2017) contend, “the cryptomarkets’ innovation originates not in the development of a new stealth technology but rather from the combination of many technologies that, when combined, provide an enhanced level of anonymity to participants” (107). In this regard, there are two key encryption technologies leveraged by cryptomarkets to ensure functional efficiency and fluid communication among participants: Tor and cryptocurrencies.

Tor (The Onion Router) is a free “circuit based low-latency communication service” which allows users to engage on the internet without revealing their location or identity (Dingledine, 2004; Mathewson et al., 2004). It is, moreover, a network within which users can search for and host an illicit website. This is particularly useful for individuals seeking to both set up a cryptomarket and conceal their hosting location from law enforcement and other aggrieved parties. Launched in 2002, TOR was initially designed by the Centre for High Assurance Computer Systems at the U.S. Naval Research Laboratory for the purposes of protecting the anonymity of government employees (Bradbury, 2014). However, as with most software designed by state actors, TOR was designed for use by state actors but trickled down to citizens once the technology was made public.

Tor uses a concept called onion routing which directs a user’s IP address through a series of random relay points to obfuscate the user’s point of origin (Bradbury, 2014, 14). “The sender of a piece of traffic will find an entry point and choose a random routing path through a selection of relays to obfuscate their point of origin. Traffic routed along this path will be encrypted until it leaves the last relay, to be sent to a specific IP address on the public Internet” (Bradbury, 2014, 12). In short, onion routing is premised on separating where you are in the world from where you are connecting to on the network (Lewman, 2016, 16). This technology is publicly accessible and easy to use. The Tor network can be accessed by using the TOR browser which is a standard web browser much like Internet Explorer, Google Chrome, or Mozilla Firefox.

Cryptocurrencies are the second major encryption technology employed by cryptomarkets. This electronic currency system allows for direct and anonymous peer-to-peer transactions without involvement or oversight from a third-party organization. Unlike fiat currencies, cryptocurrencies possess a decentralized ledger that records all transactions that have been facilitated by that respective currency (Cox, 2016). This is called the “block chain”. With the block chain, one can easily see which users hold what amount of cryptocurrencies in their digital wallet. To elaborate, a “block” is a series of updates of transfers between users. Importantly, because the block chain is a decentralized program with copies housed on all computers across the planet, transactions made with cryptocurrencies cannot be reversed, frozen, or tampered with by third-party institutions like banks and

governments. As this ledger cannot be controlled by a single entity, law enforcement cannot intervene in halting or reversing illicit financial transactions. Nevertheless, this ledger contains information about transactions which can pose a risk to those involved in illicit activity. To conceal their identity and potentially avoid prosecution if embroiled in criminal activities, a cryptocurrency user may separate their transactions from their identity.

Importantly, two studies have uncovered a strong relationship between bitcoin transactions and purchases on cryptomarkets. Janze (2017), utilizing panel data of 296,875 cryptomarket product listings as well as Bitcoin blockchain transactions, found a curious co-evolution between Bitcoin and cryptomarkets. That is, transactions within the Bitcoin blockchain and the usage of transaction obfuscation services could be reliably linked to previous sales on cryptomarkets. The author demonstrated that for one additional item sold on darknet markets, additional transactions increased by 0.123 on the blockchain six days later. As well, Foley, Karlsen, and Putnins (2018) estimated that 46% of bitcoin transactions (\$76 billion) were tied to illegal activities, many of which occurred on cryptomarkets.

Given the mandatory use of TOR (or other networks, e.g. I2P, Zeronet, Freenet, Openbazaar, etc.) and cryptocurrencies, participation in cryptomarkets requires a certain level of technical sophistication. Prospective cryptomarket participants must have a working knowledge of these technologies (Christin, 2013). Conducting qualitative interviews on Silk Road, Van Hout and Bingham (2013a) observed that training and experience with computer systems were viewed by vendors as an important skill to have. One participant noted, “If you are not a computer scientist, a lot is down to just faith. A seller has to learn a lot about the technology, if they are concerned with staying safe. It’s a big subject to dive into and much deeper than what you may initially think” (van Hout and Bingham, 2013b, 54). Aside from technological know-how, participants will also need access to a number of devices, programs, and information. In a detailed distillation of the cryptomarket literature on drug purchasing, Barratt and Aldridge (2016) observe that “prospective participants will require: a computer or equivalent device, a special anonymising browser, the marketplace URL, some cryptocurrency, a vendor willing to send the drugs to your location, and an address where the package containing the drugs can be sent” (4).

Nevertheless, anonymity networks and cryptocurrencies have created a relatively anonymous transacting infrastructure that is both opened and closed to the general public. According to May and Hough (2004, 550-551), whereas “open markets are those that are open to any buyer, with no requirement for prior introduction to the seller, and few barriers to access, closed markets are ones in which sellers and buyers will only do business together if they know and trust each other, or if a third party vouches for them.” In this regard, cryptomarkets are open to all with sufficient knowledge of anonymity networks and cryptocurrencies but is, for all intents and purposes, closed to those incapable of building rapport with customers or conducting themselves appropriately on these platforms (Aldridge and Decary-Hetu 2016; Christin, 2013, Duxbury and Haynie, 2017; Paquet-Clouston, Décarry-Héту, and Morselli, 2018).

Christin (2013) and Martin (2013) maintain that law enforcement organizations typically have more organizational experience and expertise in prosecuting terrestrial forms of illicit exchange. Moreover, the use of Tor and cryptocurrencies adds an additional layer of difficulty. It is hypothesized that the complexity of encryption algorithms that allow cryptomarkets to operate is such that it would require crackers tens of thousands of years to decrypt (Martin, 2014a, 357). Furthermore, the privacy of cryptomarkets allows for the formation of various communicative norms which are practiced in a reduced capacity in

traditional organized crime settings. This will be discussed in a later section. Importantly, it is unlikely that cryptomarkets usage will be mainstreamed. Based on 350 hours of unstructured observation during an ethnographic study, Kowalskia, Hooker, and Barratt (2019) concluded that the current levels of complexity and obfuscation constructed in the cryptomarket environment act as a barrier to the widespread acceptance of this technology. Nevertheless, as cryptomarkets continue to develop the ease of use of these platforms are bound to change and with them the likelihood that cryptomarket usage will increase.

Cryptomarkets are a remarkable criminal innovation. They provide sellers with a virtual location to advertise and sell their products to a worldwide market without constant fear of law enforcement intervention. Within terrestrial markets, this is an extremely difficult undertaking as secrecy and anonymity must be maintained by fallible human actors. Moreover, law enforcement actors are on a relatively even playing field with terrestrial criminals, able to infiltrate criminal organizations and gather intelligence without the having to deal with technological barriers. This is not to suggest that crime prevention in an offline setting is an easy task as it is marred with a bevy of other challenges unique to this environment.

Organizational Structure and Governance within Cryptomarkets

Regardless of its legality, the objectives and operational capacity of a business are often dictated by its organizational structure. The organized crime literature has long entertained discussions surrounding the horizontal or vertical composition of illicit entities. Indeed, it was once theorized by organized crime scholars (von Lampe, 2016) that the global drug trade consisted of a series of hierarchical bureaucracies that actively engaged in micromanagement and vertical integration. Such sentiments were strengthened by the media's characterization of drug traffickers in the Columbian municipalities of Medellin and Cali as cartels which restricted market competition and regulated international drug pricing (Kenney, 2007, 233). These suppositions have not been supported by Kenney (2007) and Malm and Bichler (2011) who have documented the existence of a decentralized organizational schema within these drug trafficking organizations. Nevertheless, a portion of the cryptomarket literature is dedicated to the organizational structure and associated divisions of labour within these illicit online marketplaces. In this regard, James Martin (2014a; 2014b) has been particularly instrumental in matters relating to organizational structure and governance in cryptomarkets.

According to Martin (2014a), cryptomarkets are hierarchically structured. This is reflected in some of the organized crime literature. Donald Cressey (1967) surmised that organized crime (at least in the American context) resembled an octopus, possessing one head with many tentacles. Cressey's observation would give rise to a subset of the organized crime literature which details the vertical structure of criminal organizations. Von Lampe (2016) explains that the various components within these hierarchies perform tasks which are coordinated by a common manager (105). Moreover, these vertically-structured organizations "typically have defined boundaries and internal divisions and a centralized chain of command" (von Lampe, 2016, 105). According to Catino, the basic operational unit within the Sicilian Mafia is the "family", a criminal group which possesses a territorial base from which it manages a zone or inhabited area (2014, 188). Moreover, the Sicilian Mafia possesses a vertical micro-organizational structure as these families are arranged

hierarchically with subdivisions of power which further correlate to divisions of labour (Catino, 2014, 188). To elaborate, this chain of command consists of a base of “button men” or “soldiers” who carry out operational orders, capidecina who oversee a platoon of soldiers, and a democratically elected “representative” who functions as the boss of the family (Catino, 2014; von Lompe, 2016).

This hierarchical chain of command ensures that no illegal activity occurs without permission from the boss. Importantly, however, the Sicilian Mafia also possesses a vertical macro-organizational structure. That is to say, a group of representatives will nominate a district boss who then serves as a member of a provincial commission which will collectively nominate a provincial representative who functions as a secretary and coordinator for that specific province (Catino, 2014, 190). These supra-local and provincial configurations allow for an advanced capacity in engaging the state. This, of course, corresponds with the Sicilian Mafia’s record for state-based violence. Von Lompe maintains that American Mafiosi, identical to their Sicilian brethren, have organized themselves into a number of local micro-units, the aforementioned “families” (2016, 189). As such, the American Cosa Nostra possesses a vertical micro-organizational structure with a chain of command consisting of a boss, an underboss and consiglieri, and a series of capos who supervise groups of soldiers (von Lompe, 2016, 188). Furthermore, the American Cosa Nostra possesses a vertical macro-organizational structure as the bosses of the individual families have, at one point in time, cooperated to form The Commission. The Commission functioned as a forum where the various bosses “come to agreement on matters of general importance, such as the admission of new members and the resolution of conflicts within and between the families” (von Lompe, 2016, 189). These structural orientations ensure that associates and members seek approval from their bosses and the bosses from The Commission before carrying out a criminal act.

The organizational structure of cryptomarkets is similar in some respects. From detailed and readily available sources of information, Martin (2014a) was able to identify four unique user types within the cryptomarket hierarchy. These include administrators, moderators, vendors, and consumers. Reflective of von Lampe’s (2016) observation, cryptomarkets are refreshingly transparent in their organizational structure and divisions of labour as each user type possesses a different but inter-reliant portfolio of responsibilities and capabilities (Martin, 2014a, 17). This, moreover, ensures that the platform operates fluidly as each category of actor possesses a specific role which they play. This, however, is not reflective of a purposive division of labour where each category of actor is allocated duties and responsibilities. Instead, actors will organically play a specific role by virtue of their orientation within the market’s organizational structure. This, however, is not necessarily the case for moderators as their duties and responsibilities are established by administrators.

The organizational structure of cryptomarkets is dominated by an administrative unit which oversees the efficient operation of these platforms. It is the role of administrators to act as executives, managing their site and determining the policies under which users will operate. As Martin (2014a) maintains, administrators are responsible for “authorizing and suspending individual accounts, overseeing ‘stealth’ transactions not publicly listed, creating new product categories, authorizing or prohibiting the sale of various items, as well as innovating and implementing new security procedures and cyber-defences” (18). Of course, a vital secondary function includes the management of cryptocurrency transactions. This involves the provision of escrow services which yields sales commissions from each transaction. Christin (2013), in an analysis of financial trends on the dark web, found that sales commissions typically varied between 3% and 8% of the total transaction cost.

Interestingly, Martin also suggests that administrators serve as organizational figureheads, actively engaging with media outlets and scholars (2014a, 18). Below administrators are moderators. Given their reduced administrative access, moderators assist administrators in site maintenance and customer support. This mainly involves “regulating forum discussions, identifying fraudulent activity committed by scammers, and responding to requests for assistance and complaints from vendors and consumers” (Martin, 2014a, 18). As such, moderators will perform the day-to-day activities pursuant to cryptomarket management, ensuring that the fluid operation of the platform.

At the bottom of this hierarchy are vendors and buyers. In order to operate on a cryptomarket, vendors must pay a registration fee to the site administrator. In contrast, buyers must set up a free account on the platform. From there, a vendor can set up their account and begin listing products that they wish to sell to buyers. When a buyer completes a transaction with a vendor, they have the option of providing public feedback of their experience with the vendor. This is a particularly important task as it assists vendors in building their reputation while signalling to other buyers the quality of the vendor. While vendors and buyers do not possess a formal administrative role that has been allocated to them, they nevertheless perform an essential function: engaging in voluntary economic transactions. They are, moreover, involved in two key community-building activities: product testing and friendly forum discussion. As such, these actors are responsible for the evaluation of experiential goods and creation and maintenance of a collegial community within a cryptomarket.

Nevertheless, the notion that cryptomarkets are strictly hierarchical is disputed. Norgaard, Walbert, and Hardy (2018), in an analysis of the determinants of network structure and hierarchy in physical and dark web drug markets, found that illicit online markets were generally less hierarchically structured relative to their more monopolistic terrestrial counterparts. Utilizing agent-based modelling to compare the density and average path length of their simulated black market networks, the authors found that lower transaction costs and information asymmetries in cryptomarkets resulted in less of a top-down schema compared terrestrial markets.

To this extent, Martin, in contravention to his previous findings, contends that the successful operation of a cryptomarket is deeply dependent on a decentralized exchange network between vendors and consumers (2014b, 363). Drawing on Natarajan (2006), Kenney (2007), and Bright et al. (2014), Martin (2014a) emphasizes the importance of structured economic relationships as developed through community engagement in product assessment and forum discourse. These activities serve to routinize voluntary economic transactions among buyers and vendors by developing a social fabric based on mutual interests. Brenner (2010) speculated that the various functionalities afforded by Internet connectivity would likely replace traditional organized crime hierarchies with decentralized networks. While Brenner’s contention is partially disproven by the existence of administrative units within cryptomarkets, it remains largely accurate with regard to the structure of buyer-vendor relations therein.

Terrestrial markets are often expansive international entities, comprised of many distinct individuals and groups which operate a section of the supply line (Salt and Stein, 2002; Vayrynen, 2003). This is partly the result of logistical necessity. Consider the flexible exchange networks of the Medellin and Cali cartels. According to Kenney (2007), the drugs sold by these organizations were first produced by local farmers, procured by purchasing groups, refined in specialized processing labs, exported to various trans-shipment points in the Southern hemisphere, then sold by distribution groups in a multitude of overseas markets (424). Of course, these node linkages were established by an array of brokers who introduced

interested parties and maintained lines of communication. According to Martin (2014a), cryptomarkets possess a similar exchange structure as illicit goods are created by producers, acquired by vendors (who may themselves be producers), then shipped to consumers using conventional postal services. However, cryptomarkets exchange networks differ structurally from terrestrial illicit distribution networks as they require less nodes to function (Barratt and Aldridge, 2016; Christin, 2013, Martin, 2014a, 55).

Specifically, the various brokerage services offered by cryptomarkets (escrow, direct communication, etc.) encourages vendors and buyers to adopt a direct business-to-consumer schema which serves to eliminate the involvement of drug traffickers, wholesalers, secondary brokers, and a bevy of other specialized intermediaries. As Martin notes, “Unlike conventional distribution networks, where a wide range of nodes specialise in different stages of distribution (e.g. trafficking, wholesaling), networks facilitated online are able to connect nodes and end consumers, in the absence of geographic proximity or interpersonal contact” (2014a, 55). Christin (2013), in a study of illicit online markets, observes that such innovations in drug distribution results in price stability, increased product purity, and higher levels of customer satisfaction.

Though the organizational structure of cryptomarket communities increases the efficiency of distribution networks, it is a rather remarkable development in matters relating to automation in illicit markets. Unlike terrestrial markets where interpersonal relationships are forged through friendship, kinship ties, or other personal contacts, cryptomarkets broker transactions by providing transparent and quantifiable information. As Martin (2014a) suggests, “the automation and user involvement associated with these processes mean that cryptomarkets are able to act as a kind of ‘super broker’” (45).

Importantly, monetary exchanges, brokerage services, and peer-to-peer communications are facilitated by cryptomarket administrators and the online platform itself. These illicit exchange networks are particularly durable as the elimination of a node due to arrest or competitive violence often means that the line of distribution is rerouted through an adjacent node (Martin, 2014b, 363). What’s more, the criminological literature demonstrates that familial and associational ties allow exchange networks to function as the inherent social connection and economic interdependence among participants fosters communication and market harmonization. Of course, cryptomarket vendors and buyers are not bound by familial ligatures as Italian Mafiosi or Columbian drug traffickers are (Decary-Hetu, 2016). Rather, interpersonal relations hinge on mutual interests in libertarianism, recreational drug use, and other subcultural niches (Martin, 2014a; Munksgard and Demant, 2016). In fact, the original Silk Road was founded on libertarianism (Maddox et al., 2016). Still, many participants on cryptomarkets may not abide by or even support this ideology. As such, their allegiance may lie to a trusted vendor. We will examine the topic of trust on cryptomarkets in a forthcoming section.

Importantly, a small but growing subset of the cryptomarket literature is dedicated to governance. By virtue of their illegality, environments which are conducive to the commission of crime are often perceived as spheres of lawlessness. Cryptomarkets, in particular, engender additional considerations as they are almost bereft of any external regulation or government oversight. “Products that are sold on illicit sites bypass the processes of government-mandated testing, quality control, and safety standardisation that are imposed on regular consumer goods” (Martin, 2014a, 37). However, cryptomarkets are, in actuality, more collegial than other illicit organizations. According to Martin (2014a), “the synergies produced by mentalities, technologies, resources and institutions allow cryptomarkets to function as sites of informal nodal governance” (18). That is to say, these

illicit online marketplaces possess clearly defined rules that are developed and enforced by administrators.

The objectives of a business are often dictated by its organizational structure. Catino stipulates that vertically-structured crime organizations are “characterized by the presence of higher levels of coordination, centralized power, and systemic decision-making processes” (2014, 177). In this regard, Martin (2014a) maintains that the centralization of power within the hands of cryptomarket administrators encourages nodal governance. Originally conceived by Shearing and Wood (2003), nodal governance suggests that in the absence of formal government, non-government, and commercial institutions, informal groups emerge as substitute pseudo-governments. These informal groups carry out a range of regulatory functions including contract and rule enforcement, dispute resolution, security, and policing.

Indeed, nodal governance is rife among various criminal organizations. Diego Gambetta (1993) contends that the Sicilian Mafia arose in the 19th century following Italy’s shift from feudalism to capitalism. Specifically, the emergence of burgeoning markets resulted in predatory attacks from which the Italian government could not provide protection. As such, the gabbellotti (the precursor to the Sicilian Mafia) enforced private property rights whilst providing impromptu governance. Though illegal markets are assumed to be stateless entities which lack formal conflict resolution mechanisms, the scholarly literature suggests otherwise (Reuter, 1985). Indeed, order can spring organically from iterated engagements from actors operating in an illicit environment. These repeated interactions create norms and customs that proliferate among the criminal actors. If these norms and customs are not followed by actors in future interactions, noncompliant actors will be punished compliant actors. This can vary from naming and shaming to ostracism.

As it pertains to cryptomarkets, the absence of a formal regulatory body has created an environment ideal for the guiding hand of administrators. In a study of conflict resolution mechanisms in cryptomarkets, Morselli et al. (2007) observed the proliferation of official rules which governed the conduct of buyers and vendors. Established by administrators, these rules often took two forms: moral/ethical and functional. Moral/ethical rules banned the sale and distribution of particular items (child pornography, firearms, etc.) while functional rules prevented thefts and scams which might undermine market efficiency and interpersonal trust. Consider Ross Ulbricht’s installation of the Silk Road Charter, a utopian constitution of sorts which guided user interaction (Martin, 2014a, 13). The Silk Road Charter described the Silk Road as a “global enterprise” whose guiding principles revolved self-ownership, personal responsibility, user equality, personal integrity, and a commitment to self and communal improvement (Martin, 2014a, 12).

Furthermore, cryptomarket administrators play a pivotal role in both enforcing marketplace rules and adjudicating disputes between buyers and vendors (Martin, 2014a; 2014b). Norgaard et al. (2018) documented the dispute resolution procedures of Hansa, the third-largest cryptomarket in 2017. Interestingly, the authors observed that Hansa encouraged buyers and vendors to solve disputes among themselves but would resolve the dispute if a private resolution was not possible. In fact, if the dispute is ruled in favour of the buyer, Hansa administrators would force the vendor to compensate the wronged buyer.

With relation to nodal governance, Zajacz (2017) contends that cryptomarket administrators took on “a key function of the state: protecting citizens from harming each other through force, fraud, or theft” (78). To this extent, a participant in Van Hout and Bingham’s study of Silk Road vendors noted that “We are a community, and Dread Pirate Roberts (Silk Road administrator) is our president in a sense” (2013a, 35). Importantly, the

mode of governance assumed by administrators somewhat reflects Varese's (2010) definition of an organized crime group (OCG). Varese defines OCGs as a "group which attempts to regulate and control the production and distribution of a given commodity or service unlawfully" (2010, 14). In this regard, Martin (2014a) notes that administrators engage in creating new product categories, authorizing or prohibiting the sale of items, and overseeing all transactions (18). As cryptomarket administrators manage their platform's escrow service, they possess complete oversight over all formal transactions conducted between vendors and buyers. Indeed, administrators approved cryptocurrency transfers to vendors upon receiving confirmation of product delivery from consumers, receiving a sales commission in the process (Martin, 2014a, 31).

Indeed, Cryptomarkets, aided by their hierarchical structure, further resemble Mafia groups by assuming a protective role, presiding over all darkmarket transactions. To elaborate, as cryptomarket administrators manage all cryptocurrency transactions the implication is that they possess complete oversight over all interactions between vendors and consumers. Indeed, administrators, asserting their protective authority, may approve a currency transfer upon receiving confirmation of product delivery from the consumer (Martin, 2014a, 31). Similar to Mafias, administrators may punish those who renege on their contractual obligations (i.e. dark market fraudsters) by suspending their account. Furthermore, Mafiosi have the important task of protecting their clients from law enforcement (Varese, 2010, 17). Of course, Administrators must actively shield vendors and consumers from law enforcement by "innovating and implementing new security procedures and cyberdefences" (Martin, 2014a, 18). Administrators, functioning as "capable guardians", prevent marketplace actors from being defrauded (Cohen & Felson, 1979). Furthermore, inbuilt conflict reduction mechanisms create a fairly well-regulated market.

However, it is important to acknowledge that cryptomarket administrators, unlike Mafiosi, do not engage in extortion and are rather benevolent in thought and action. In summation, the hierarchical structure of the administrative unit in cryptomarkets is predicated on nodal governance which pertains to the simultaneous regulation of illicit commodities and provision of protection. Still, this is not to suggest that the motivations of cryptomarket administrators and mafiosi are the even remotely similar. This comparison between the two is done as a means of demonstrating that cryptomarkets possess some elements which resemble a traditional organized crime group. Furthermore, this does not mean that cryptomarkets are themselves organized crime groups or that governance therein is an uncomplicated process. Those who suggest that illicit online markets are organized crime groups often point to the hierarchical structure of these entities as evidence. However, as Lusthaus (2012) suggests, the provision of a secure space for illicit transactions, restriction of access to deviant members, and third-party enforcement of contracts are not themselves qualities which make an illicit online marketplace an organized crime group. Indeed, a mere structural design does not make an illicit entity an organized crime group.

Furthermore, Lusthaus (2012) offers three reasons for the ineligibility of illicit online markets as organized crime groups. First, illicit online markets are individual marketplaces rather than regulatory bodies which preside over entire industries. Indeed, while the Sicilian mafia did not itself sell fish it was firmly in control of Palermo's fish market. To this extent, cryptomarkets are merely brokerage platforms which allow vendors to advertise their wares and buyers to purchase them. While administrators can ban the sale of specific goods and services (e.g. child pornography) they have no control over the supply and demand of products which are advertised on their platform. Secondly, online markets are comprised of autonomous groups and individuals which lack a single, coherent objective. Intuitively,

cryptomarkets participants have no overarching objective outside of engaging in voluntary economic transactions. There are no broad organizational goals which tied individuals together.

Finally, governance, namely protection and the enforcement of contracts, is especially difficult to actuate as violence, a key regulatory tool for terrestrial organized crime groups, cannot be exercised in cyberspace. While violence is a constitutive element in regulating and governing illicit markets, its employment comes at tremendous costs as it draws unwanted attention from law enforcement, normalizes violent retaliations, and discourages potential customers and partners (Reuter, 1985; Campana and Varese, 2013). Open markets and illicit markets whose lack of barriers to entry permits for unregulated admissions, are privy to violent turf wars as dealer-dominated locales are subject to unwanted incursions from new competitors. Nevertheless, the anonymity and geographical dispersion afforded to cryptomarkets means that participants cannot simply harm other disreputable actors. The improbability of violence in cryptomarkets lies in the platform's dematerialization of voluntary economic transactions. Indeed, the cryptomarket literature is replete with studies that point to broad-based reductions in violence among users (Aldridge and Decary-Hetu, 2014; Morselli et al., 2017; Van Hout and Bingham, 2013a).

Mohamed and Fritsvold (2010) found that cryptomarket vendors have a reduced likelihood of violence enacted against them compared to "street" dealers as most of their clientele are middle-class, university students who are generally averse to serious interpersonal violence. Furthermore, Barratt et al. (2016), surveying 3794 respondents from 57 countries on drug use, found that 1.3% and 1% of cryptomarkets users experienced "threats to personal safety" and "physical violence", respectively. In contrast, 14% and 6% of those who purchased from friends, 24% and 10% of those who purchased from known dealers, and 35% and 15% of those who purchased from strangers experienced "threats to personal safety" and "physical violence", respectively. The authors concluded that "cryptomarkets are associated with substantially less threats and violence than terrestrial market that are also used by cryptomarket customers" (Barratt et al., 2016, 20). In general, buyers reported safer and more convenient transactions given the complete circumvention of face-to-face meetings with potentially dangerous dealers (Barratt, Lenton, and Allen, 2016; Van Hout & Bingham, 2013a, 2013b).

Still, the relative absence of animosity much less violence within cryptomarkets may be a by-product of vendor behaviour. As Martin indicates, the cryptomarket vendors are encouraged to create a "socially constructive public image that is both free from violence and more attuned to the perceived priorities of their customer base" (2014a, 40). Moreover, there is a certain futility to violence in cryptomarkets as it retains no strategic value. Indeed, financial success of a cryptomarket vendor is often contingent on more benevolent qualities. This is substantiated by Aldridge and Decary-Hetu (2014) who state that "a different set of skills is required of cryptomarket vendors to succeed (e.g., good customer service, writing skills) compared with conventional dealers who can utilize physical intimidation to maintain market share" (25). To this extent, creating rapport and behaving in a trustworthy manner go farther on cryptomarkets than would violence were it a readily available option.

This is not to suggest that the inability to use violence precludes any malicious activities that may hamper the operation of a cryptomarket. Buyers and vendors are often victims of scams perpetrated by other participants. As Morselli et al. (2017) suggest, "the most common scams are thefts by vendors (when lying about having shipped the drugs) and buyers (when lying about not having received the drugs)." Furthermore, bad management by cryptomarket administrators is often singled out by patrons as the primary reason for why

scams are allowed to take place (Martin, 2014). This typically comes in the form of negligence where administrators are slow to take action against bad actors on their platform. This mismanagement by cryptomarket administrators reduces consumer confidence and creates instability within a market. However, this inability to wholly govern the conduct of actors on a cryptomarket is perfectly understandable given the level of anonymity and encryption on these platforms. As such, a more formal top-down form of governance is particularly difficult to achieve on cryptomarkets.

Nevertheless, despite the difficulties of top-down governance in cryptomarkets, buyers and vendors can resolve conflicts amongst themselves. In *Disorganized Crime*, Peter Reuter (1985) argues that while illegal drug markets are “stateless” entities, participants can themselves resolve their disputes without aid from a regulatory body. In this regard, Morselli et al. (2017), in an investigation of several cryptomarket forums, found six peer-to-peer conflict resolution strategies that are traditionally absent in off-line drug markets. These include: 1) demonstrating tolerance/patience when dissatisfied, 2) avoiding conflict and refusing to intervene, 3) ostracizing or naming and shaming bad actors, 4) levying threats against bad actors, 5) negotiating a private settlement, and 6) calling upon a third-party mediator. Furthermore, Morselli et al. (2017) note that four channels are available to cryptomarket participants when conflicts arise. Participants can initiate direct contact through built-in messaging systems, use the formal support ticket system to notify administrators of a conflict, publicly shame disreputable actors on a forum, or damage the vendor’s reputation by leaving negative feedback.

Governance on cryptomarkets is characterized by an innovative combination of private ordering and nodal governance. This hybridity has various functional advantages including greater ease of operation, transparent communication, and greater awareness of consumer satisfaction. In general, a decentralized exchange network increases the fluidity of voluntary economic transactions while a competent administrative unit supervises these transactions to ensure satisfaction among all parties involved in a transaction. This mixture of governmental paradigms aids in the function of these illicit online markets. In this respect, cryptomarkets are unlike traditional organized crime groups as they are both marketplaces and mediators with enhanced communication and anonymity. Moreover, it is this encryption and anonymity which reduces the capacity for a more traditional top-down form of governance on cryptomarkets.

The organizational structure of cryptomarkets can be described as it is an innovative stitching of a hierarchical administrative unit and a decentralized exchange network. Of course, such hybridity is rarely documented in the organized crime literature as it seems illicit entities are primarily horizontal or vertical in structure. Moreover, a cryptomarket’s amalgamation of organizational features from Mafias and drug trafficking organizations is a rarity in and of itself. Nevertheless, this hybridity has various functional advantages. A decentralized exchange network increases both the durability of the distributive chain and the fluidity of voluntary economic transactions while a powerful administrative unit oversees the legitimacy of these transactions while establishing codes of conduct. This mixture of precise governance and free market economics is appropriate in illicit markets. Indeed, the lack of conventional enforcement mechanisms in illicit markets necessitates the existence of an overarching entity to moderate market transactions and punish those who renege on contractual obligations. Nevertheless, the supremacy of cryptomarket administrators often means that their elimination may topple the platform. On the other hand, the loose network of vendors and consumers allows for user mobility and the expedient rebirth of the platform.

The Who, the What, and the Where of Cryptomarket Studies

The vast bulk of scholarly literature on cryptomarkets are either descriptive or qualitative (Baratt and Aldridge, 2016). To elaborate, descriptive studies document the range, type, and quantity of illegal goods and services (Aldridge & Decary-Hetu, 2014; Martin, 2013a) while qualitative studies seek to identify the characteristics and motives of cryptomarket participants through interviews with buyers and vendors. In this respect, these studies can be neatly separated into three categories of query and investigation: what items are sold, who sells them, and where they are shipped to/from. It is important to stipulate that this is not an exhaustive categorization as there are several studies which do not fall into any of these categories. These include studies of trust, confliction resolution, and network structure.

Reflecting its wide array of illicit wares, the motto of Silk Road was “If you can smoke or, inject it, or snort it, there’s a good chance Silk Road has it” (Goodman, 2016). Indeed, cryptomarkets offer a prodigious selection of drugs, malware, weapons, credit card and banking information, airplane tickets, counterfeit money, child pornography, chemical substances, and hitmen, among other illicit goods and services (Baratt and Aldridge, 2016; Christin, 2013; Decary-Hetu, Mousseau, and Rguioui, 2017; Hutchings and Holt, 2015; Martin 2014a). Although it is difficult to quantify the exact number of transactions conducted on cryptomarkets, Aldridge and Decary-Hetu (2016) have developed a suitable metric: counting the number of buyer feedback messages on a listing. While not perfect, this method provides researchers with reliably accurate estimates of monthly revenue by item type.

Most to-date studies of items sold on cryptomarkets have fixated on controlled substances. In an extensive study of 16 marketplaces, van Buskirk, Naicker, Roxburgh, Bruno, and Burns, Breen, and Roxburgh (2016), identified “cannabis, pharmaceuticals, MDMA, cocaine and methamphetamine as the five most commonly sold substances, with the popularity of new psychoactive substances declining slightly” (20). With regard to the sale of fentanyl, fentanyl analogs, and other synthetic opioids, researchers (Lamy et. al, 2020) found that DreamMarket, the largest cryptomarket in history, offered a steady supply of synthetic opioids at both retail and whole-sale prices. Curiously, China was the main country of origin of novel synthetic opioids while 52.6 % of all fentanyl-type drug listings were posted by unique vendor names who indicated they were shipping from the U.S. and Canada.

Utilizing digital trace to examine the prevalence of nonmedical prescription psychiatric drug use on 31 cryptomarkets, Cunliffe, Decary-Hetu, and Pollak (2019) found that diazepam, alprazolam, Adderall, modafinil and methylphenidate were the most popular sedatives and CNS stimulants. Moreover, the US and UK were the primary suppliers of these products, accounting for 41.4% and 31.1% of all sales, respectively. Surprisingly, antidepressants and mood stabilisers were not particular popular on cryptomarkets. The authors conclude that only the nonmedical prescription psychiatric drugs that have a potential for abuse are sold at high levels.

As it relates to the popularity drugs on cryptomarkets, there appears to be a remarkable level of consistency from market to market. “Since 2015, cannabis, MDMA (ecstasy) and cocaine-related products have been the most popular drugs sold online, representing about 70% of all sales” (Soska & Christin, 2015, 55). This consistency seems to suggest that cryptomarkets cater to specific drug types over others. Nevertheless, the size and scope of smaller niche drug markets on the dark web has increased in recent years. In a study of the six largest cryptomarkets, Barrera, Malm, Decary-Hetu, and Munksgaard (2019) found that tobacco sales reached US \$194,940 annually as a lower-bound estimate. Of importance is Barratt et al.’s (2016) Global Drug Survey (N=3794) which found that MDMA/Ecstasy

(55%), cannabis (43%) and LSD (35%) were the drug types most commonly obtained through cryptomarkets.

Nevertheless, a novel subset of these item-centric studies seeks to uncover the effect of changing drug policies on both cryptomarket sales and rates of user harm in light of the growth of these marketplaces. Using an interrupted time series analysis, Martin, Cunliffe, Decary-Hetu, and Aldridge (2018) investigate the association between the rescheduling of hydrocodone products in the US and the subsequent increase of illicit prescription opioids on cryptomarkets. The authors found that the opioid market share on cryptomarkets increased from 6.7% to 13.7% following the implementation of the hydrocodone rescheduling. Moreover, there was a statistically significant change in the composition of the opioid market as fentanyl sales spiked dramatically. However, despite increases in the use of harmful drugs, cryptomarkets may also decrease the deleterious effects of drug use by providing harm-reduction information to buyers. To this extent, the benefit of these analyses is their elucidation of consumer preference for cryptomarket transactions (van Buskirk, Roxburgh et al., 2016). Indeed, drug quality seems to be a major factor in consumers' decision to use cryptomarkets. Caudevilla et al. (2016) reported results of laboratory testing of samples sent by cryptomarket vendors. In general, the authors found that the samples were purer and less adulterated relative to samples provided by terrestrial sources.

Much of the early cryptomarket literature had an explicit focus on either analysing Silk Road data or examining user experiences and vendor characteristics through interviews (Maddox et al., 2016). As it pertains to qualitative assessments, studies by Van Hout and Bingham (2013a, 2013b) examined consumers' decision-making processes and motivation for participating on cryptomarkets. Relying on a single case study, van Hout and Bingham (2013a) insisted that the variety of controlled substances mixed with cryptomarket reputation dynamics encouraged user participation. More importantly, the same authors (van Hout and Bingham, 2014b), monitoring discussion threads and conducting anonymous online interviews (N = 20), found that users frequented the Silk Road out of curiosity and concerns for personal. This study was the first to examine the demographical breakdown of cryptomarket participants; reporting that the majority of users were white males between the ages of 18 to 25 who preferred MDMA, ketamine, cannabis, and cocaine. None of the findings from this study can necessarily be generalized due to the small sample size. Three additional papers (Bancroft and Reid, 2016; van Buskirk, Roxburgh et al., 2016) corroborate these findings.

As it pertains to consumer participation, however, Barratt, Lenton, Maddox, and Allen (2016) conducted a digital ethnography which spanned two years (2012-2014) and included 17 Silk Road buyers. These in depth and unstructured interviews revealed that consumer participation on cryptomarkets amounted to being "a kid in a candy store" with high product availability reducing the need to hoard drugs, and by extension, helping to moderate drug use. The honeymoon period that is often experienced by buyers upon successfully purchasing drugs from a cryptomarket for the first time transforms into a stable or decreasing trajectory of drug use (Bancroft and Reid, 2016). Similarly, Van Hout and Hearne (2016), examining cryptomarket forum members' views and perspectives on new psychoactive substances (NPS), found that buyers "appeared well informed, with harm reduction and vendor information exchange central to purchase decisions". In general, Van Buskirk, Roxburgh et al. (2016) have reported that cryptomarket consumers are typically a more "entrenched" consumer group with active ties within their own community.

Studies by Aldridge and Decary-Hetu (2016b), Decary-Hetu, Mousseau, and Vidal (2018), and Demant, Munksgaard, and Houborg (2018) demonstrate that cryptomarkets increasingly cater to business-to-business transactions and social drug dealing as opposed to simple business-to-consumer transactions. In this regard, the potential for the dark web's globalization of the drug trade is demonstrated by several studies. With regard to the geographical distribution of cryptomarket activity, Dolliver, Ericson, and Love (2018) found that Australia, Canada, Germany, the Netherlands, the United Kingdom, and the United States accounted for the largest number of listings and transactions for heroin, cocaine, and prescription drugs. Moreover, while heroin and cocaine are generally mass-produced in South Asia and South America, the products originated from the U.S., Australia, and the Netherlands; nations well known for consuming these drugs.

Similarly, Van Buskirk, Naicker, Roxburgh, Bruno, and Burns (2017) found that the majority of drug listings on the Agora market originated in the United States, the United Kingdom, Australia, China, and the Netherlands. These nations accounted for 61.8% of all identified listings and 68% of all unique vendors. Interestingly, Australia possesses the highest per capita estimate of sellers with 4.73 sellers per million. This makes intuitive sense as Australia's geographic isolation and relatively high drug prices encourages more of a domestic market which sells to Australian customers. Indeed, Australia is rather unique with regard to cryptomarket drug trading as studies (Barratt, Ferris, Winstock, 2014; Cunliffe, Martin, Decary-Hetu, and Aldridge, 2017; Phelps and Watt, 2014) have revealed a dense domestic market with higher than average drug prices.

A study by Broseus, Rhumorbarbe, Morelato, Staehli, and Rossy (2017) which examined the geographical structure of drug trafficking on the Evolution marketplace demonstrated that countries within the Europe and Anglosphere accounted for an outsized portion of sales and listings on Evolution. To this extent, 64% of drug listings and 30% of sales came from the U.S. Importantly, Broseus et al. (2017) also demonstrate a modicum of product specialization as niche prescription drugs were shipped primarily from the Netherlands (98% of listings), Canada (97%), Spain (96%) or Sweden (94%). While Tsuchiya and Hiramoto (2021) also found that cryptomarket transactions more often took place in Europe, the US, Canada, and Australia, transactions were more frequent on Monday, Tuesday, and Wednesday relative to Saturday and Sunday. This suggests that cryptomarket users make drug purchases between Mondays and Wednesdays for personal use on the weekend. This coheres with Aldridge and Décarry-Héту (2016), Barratt et al. (2016), and Demant et al. (2018) who maintain that cryptomarket drug purchases are recreational in nature as opposed to wholesale transactions.

Demant, Munksgaard, Decary-Hetu, and Aldridge (2018) characterized cryptomarket buyer behaviour through product reviews posted on 15 cryptomarkets. The authors found that there is an increasing movement toward the localization of cryptomarkets with regard to product destinations. Norbutas (2018), using publicly available crawls of the cryptomarket Abraxas, found that buyers were more likely to buy from multiple sellers within a single country, avoiding purchases from countries which were different. Norbutas (2018) concluded that online drug trade networks, similar to terrestrial networks, are "heavily shaped by geographic constraints in spite of their ability to provide access for end-users to large international supply" (96).

Cryptomarkets offer a wide variety of illicit goods and services. However, these products are generally bought and sold by individuals residing in developed countries across the Western hemisphere. As such, cryptomarkets are platforms utilized primarily by educated and well-to-do individuals who have a high level of technological savvy relative to the rest of

the general population. Moreover, these illicit platforms generally cater to the trade of illicit substances like marijuana, cocaine, psychedelics, and prescription drugs. However, there is evidence to suggest that newer and more dangerous drugs such as fentanyl are increasingly sold on some cryptomarkets. A number of cryptomarkets have, nevertheless, banned the sale of this product given the dangers associated with its use. The who, what, and where of cryptomarket studies constitute the early years of research on this topic. The scholarly research has evolved since then, encompassing topics such as trust, governance, and disruption. Indeed, the topics of trust and market disruption will be covered in the forthcoming sections.

Trust and Reputation

Information concerning the quality of experiential commodities is both scarce and unreliable in illicit markets. Moreover, reputations are difficult to establish and state protection is a near impossibility unless one renounces their criminal ways and turns state witness (Campana and Varese, 2013). Trust is minimal, and betrayal is the standard operating procedure. Wright and Decker (1994) and Hamill (2011) observed that betraying one's friends, family, and associates is normal in the criminal underworld. The fragility of trust within the criminal underworld may be owed to the unflattering selfishness and proclivity for risk-taking which characterizes many criminals (Gambetta, 2009, 30). Indeed, the situational constraints with which a criminal must contend (death, arrest, betrayal, etc.) certainly encourages thoughts of renegeing on contractual obligations and turning tail when circumstances dictate. To make matters worse, these contractual obligations are not upheld by a principal authority as they would be in licit markets. Nevertheless, trust is the tool which allows criminals to cooperate, ultimately permitting the heist, assassination, or arson to move forward.

Gambetta (2000) defines trust as “a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action, both before he can monitor such action (or independently of his capacity ever to be able to monitor it) and in a context in which it affects his own action” (217). Von Lampe and Johansen (2004) offer a simpler definition, noting that trust is a mechanism for individuals to “cope with risk and uncertainty in interactions with others” (103). Trust, then, involves the presupposition of future risk. It is a wager as to whether one's trading partner or accomplice will fulfil their stated or expected obligations such that the transaction or activity is successful. There are, of course, no guarantees of the fulfilment of these obligations in the criminal world.

It is, moreover, important to consider how one determines whether or not their fellow criminal can be trusted. Williamson (1993) maintains that this requires a trustee to demonstrate to the truster a temporary suspension of selfish desires for the sake of cooperation (458). The trust deficit within the criminal world is particularly problematic for trusters. As Gambetta and Bacharach (2001) demonstrate through a game theoretic approach, the optimal outcome for a trustee is to cheat (renege) in the event that a truster opts to cooperate (endow trust). However, Gambetta (2009) also contends that the iteration of this outcome over several rounds would be most detrimental to a trustee. That is, their persistent duplicitousness would discourage the truster from cooperating, costing the trustee all future business opportunities (38). Under such uncertainties and moral looseness, one can understand the fragility and paucity of trust within the criminal world.

While the criminological literature (Gambetta, 2000; von Lampe and Johansen, 2004; Gambetta, 2009; Campana and Varese, 2013) has emphasized the trust deficit within criminal

networks of varying size, these observations reflect criminal activities which take place in terrestrial markets. One would be correct in assuming that trust dynamics are liable to change in cyberspace. Several studies (Holt and Lampke, 2010; Yip, 2011) have indicated that market-driven dynamics are present in illicit online markets. Brenner stipulates that cybercriminals operate as “free-trading entrepreneurs” when engaging in drug dealing. Voluntary economic transactions are the common operational protocol. However, it is important to stipulate that transactions between cybercriminals are not predicated on thick trust as set out by Khodyakov (2007) or bonding capital as mentioned by Lo (2010). These relationships are instead built on a sort of superficial or thin trust (Khodyakov, 2007) which is easily built and destroyed as it is predicated on circumstance rather a deep connection between those involved.

Decary-Hetu and Dupont (2012), in an examination of a botnet forum, found that simple indicators often determined how well a vendor was trusted. These generally included the number awards received or size of one’s network. In this case, surface-level trust was built upon personal characteristics and behaviour as opposed to mutual experiences where a deep trust could be developed. However, another study by Dupont, Cote, Savine, and Decary-Hetu (2016) revealed that “reputation systems within botnet forums are heavily biased according to the position of the rater within the system”. That is to say, new forum members were less likely to post negative reviews or assessments. The majority of negative feedback came from forum staff and administrators. Interestingly, only a tiny fraction of the forum membership (2.4%) participated in the vast majority (75%) of “trust exchanges”.

Trust plays an important role in cryptomarkets. Lacson and Jones (2016) contend that the creators of cryptomarkets concentrated their efforts on building cohesion and camaraderie among users of various functional stripes. In researching relations between vendors and consumers, Van Hout and Bingham maintain that these relationships “were based on levels of trust and professionalism” (Van Hout & Bingham, 2013, 387). Consumers and vendors must trust one another to communicate the quality of products and the requisite currency to merit an exchange. In this case, trust is formulated through each party’s communication of their expectations when participating in a voluntary transaction. To this extent, cryptomarket users openly share information on the quality of drugs and their value relative to street-level pricing on forums. A study conducted by Dasgupta et al. (2013) found that “buyers were able to provide a valid estimate of the street price of diverted prescription opioids...and predict the relative pharmacologic potency of opioid molecules” (178). Of course, this quality of information is not easily available in conventional criminal markets as the lack of trust and need for secrecy equates to a lack of reliable information on experiential goods.

To this extent, the commission of a crime is often dependent on a criminal’s ability to obfuscate their interpersonal exchanges. That is, increasing the difficulty of law enforcement in detecting their intentions (Gambetta, 2009). Traditionally, this encrypted communication among criminals came in the form of face-to-face interactions in noisy clubs and isolated golf courses and long-distance communications with disposable mobile phones and secret radio frequencies (Gambetta, 2009, 155). Importantly, the application of encryption technology and direct messaging systems by cryptomarkets has automated secrecy and privacy.

Cryptocurrencies and anonymity networks are far more effective at befuddling law enforcement than the methods employed by other criminals. As such, cryptomarket users need not concern themselves with proactively restricting communication so as to maintain secrecy. These processes are automated through the use of routing software and cryptocurrencies. This is not to suggest the indomitability of cryptomarket encryption as human errors can always be made. Nevertheless, this efficiency in obfuscation simplifies

cryptomarkets communication as users can be forthright about their illicit intentions and desires. This can involve the establishment of consumer pricing by simply posting prices on a vendor page or voicing one's opinion on the US War of Drugs on a forum. This allows for the formation of transparent communicative norms that are practiced in a reduced capacity in terrestrial crime settings (Aldridge and Decary-Hetu, 2016a; Martin, 2014a). This bodes well for personal promotion and open communication.

Such fluidity and transparency are not present in the physical underworld given the complexity of signals and their failure to fully conceal one's criminal endeavours. As it pertains to complexity, consider the convoluted process of soliciting sex among gay men during periods when homosexuality was illegal and/or stigmatized. Gay men created "polari", "a lexicon for secret communications between gays who knew each other as gays but also a bait to check whether someone was gay and interested in making contact" (Gambetta, 2009, 166). Additionally, homosexuals developed the hanky-code which utilized a series of coloured handkerchiefs which denoted the specific sexual act one desired (Gambetta, 2009, 166). In order for polari and the hanky-code to be successful, proponents of these strategies must have a working knowledge of their operation and must moreover depend on others having similar proficiency. It is entirely possible for communication to be hampered by the complexity of these strategies.

Conversely, let us consider the simplicity of purchasing a hard-to-procure item such as uranium or hacked government data on a cryptomarket. I have purposely chosen these items as their acquisition is made difficult by dense regulation in conventional criminal markets. On a cryptomarket, transactions of this magnitude are actualized by a consumer identifying a suitable vendor, selecting the amount of the desired item, and initiating a cryptocurrency transfer. In this example, a request for uranium is made to a vendor, the vendor ships the uranium to the buyer, and an electronic currency transfer is authorized by a cryptomarket administrator once the buyer receives the product. There are no secret languages or coloured handkerchiefs. Communication is direct and transparent.

Secondly, while cryptomarket communication accurately conveys a user's message, communication in the physical underworld is handicapped by the possibility of misinterpretation. A successful signal must be disseminated throughout an organization or market, overcome variations in meaning, and mitigate the issues of memorization and ambiguity (Gambetta, 2009, 171). Given these strict hurdles of success, failed illicit business transactions are a constant throughout the underworld. The inadequacies of conventional criminal communication are such that even the most organized entities like the Sicilian Mafia have difficulty relaying messages accurately (Gambetta, 2009, 173).

A popular example pertains to a misunderstanding concerning the meaning of silence among Italian Mafiosi in the 1960's (Gambetta, 2009). Salvatore La Barbera's silence over the murder of fellow Mafiosi, Calcedonio Di Pisa, indicated to some that La Barbera had indeed murdered Di Pisa. However, La Barbera had not in actuality murdered Di Pisa as his silence was merely a reflection of his decision to remain respectful in lieu of the death of a fellow gangster. Nevertheless, La Barbera's silence was perceived as guilt, and he was assassinated (Gambetta, 2009). His death resulted in an internecine war among the various families in Palermo. Such broken lines of communication and the resultant disaster are not present in cryptomarkets as users are able to freely communicate their intentions over a transparent medium. Simply put, digital encryption has streamlined interpersonal relations in criminal markets, simplifying communications and increasing criminal extroversion as one need not worry about restricting communication to maintain secrecy.

Gambetta (2009) emphasizes that in order to successfully engage with prospective customers, sellers must themselves engage in personal promotion, solidifying their identity as a supplier and the quality of their products. Given that many illicit goods and services are experiential, a specific trademark is required to establish an association between a dealer and their product. Heroin dealers, for example, have long utilized delivery bags with unique stamps in order to sell to customers (Gambetta, 2009). These stamps delineate the dealer's identity and associated heroin product which they traffic. Similarly, cryptomarket vendors possess a customizable seller page which not only distinguishes them from other vendors but more importantly allows them to establish direct relations with buyers (Martin, 2014a, 35). To this extent, vendors adopt unique usernames to both distinguish themselves from other vendors and establish a direct relationship with customers. Some names include "Cannabis Connection" and "Dr. Leary", but may even include popular references like "haizenberg", "MrWhiteInc.", "Nancy Reagan", and "ReDEyEsEmporium" (Martin, 2014a, 35).

However, by virtue of how the dark web markets operate, names are transient and entities that are trusted today might not exist tomorrow. Vendors will try to build a reputation associated with the persona they have created. Reputations are vitally important in facilitating criminal engagements as criminals who wish to collaborate often have no prior knowledge or experience with each other and need a mantle with which to place their trust (Yip, Webber, Shadbolt, 2013, 526). Moreover, given the economic uncertainty and lack of accountability in the criminal world, a good reputation is just as if not more valuable in illicit markets than it is in licit ones (Gambetta, 2009, 199). Backed by legal and moral assurances, a reputation in a licit market may function as a coordination device which allows buyers to discern who the most trustworthy vendor is from a list of vendors selling the same products within the same price range (Przepiorka and Aksoy, 2017). This is the same in illicit markets. To this extent, Leeson (2005) makes the argument that users "need to establish ex ante whether or not the outsiders they would like to trade with are 'cheaters' or 'cooperators'" (79).

Formal institutions which collect information on parties involved in trades may not always exist in the criminal world. While not an example of an illicit exchange network, Greif (1989) documents the Maghribi traders would organize coalitions in Medieval Europe in order to exchange information about their agents' reputation to mitigate issues relating to trust amid long-distance trade. Hillmann and Aven (2011) describe a similar situation in Russia around the turn of the nineteenth century whereby the reputation of individuals was pivotal to the development of corporate capitalism. It is often the case that reputations are established and maintained through a formal system. According to Milgrom et al. (1990), the Champagne Fairs in France used bookkeeping and cashless payments as a private adjudication system that allowed them to track fraudulent traders and exclude them from future fairs. Moreover, Jappelli and Pagano (2002) maintain that credit bureaus emerged in the late 19th century, collecting and sharing information about borrowers' credit histories in order to create reputational incentives for repayment.

Nevertheless, a good reputation is difficult to develop and divulge in the criminal world. Reuter (1985) contends that the elevated risk of detection by law enforcement prevents for the willing and consistent dissemination of information among dealers, brokers, and buyers. That is to say, a consumer's unfamiliarity with a dealer's reputation for honesty or a lack thereof may be attributed to the general paucity of information circulated in a criminal market. This, however, is not necessarily the case when it comes to cryptomarkets.

On cryptomarkets, vendor reputations are established by consumers who are encouraged by administrators to provide publicly available feedback on their experience with a vendor. "Customer feedback takes a variety of forms, ranging from detailed comments

about shipping times, ‘stealth’ measures and the perceived potency of illicit drugs, to a simple 5-star rating” (Martin, 2014a, 41). Furthermore, a cryptomarket vendor cannot alter the feedback published on their page, whether positive or negative (Martin, 2014a, 42). As such, reputations cannot be artificially inflated by self-serving vendors. What’s more, these reputations are presumably up-to-date as consumers often upload feedback upon receiving their requested product (Hout and Bingham, 2013). Tzanetakis, Kamphausen, Wersé, and von Laufenberg (2016) argue that customer feedback in cryptomarket creates trust in an otherwise uncertain environment. Still, a vendor’s reputation constitutes public knowledge on a cryptomarket as it is brazenly displayed on their seller page (Christin, 2013). Vendors with a reputation for timely and trustworthy transactions have a strong incentive to behave cooperatively (Shapiro, 1983).

Hardy and Norgaard (2016) use data on cannabis listings from Silk Road to study the relationship between reputation and prices. The authors show that “reputation acts as a sufficient self-enforcement mechanism that allow transactions to occur” (Hardy and Norgaard, 2016, 32). To this extent, vendor reputations constitute a formal institution that creates a stable trading environment among those least expected to deal honestly. Janetos and Tilly (2017) show that a mature highly-rated cryptomarket vendor charges 20% higher price than a mature low-rated vendor. In general, vendors with more reviews charge a higher price than sellers with a low number of reviews regardless of rating. However, bad (i.e. low-ranked) sellers prefer to exit the market than decrease their prices in response to negative feedback. This is in line with Batikasa and Kretschmera (2018) who, studying the Agora marketplace, found that cryptomarket vendors were more likely to leave the market when they received negative feedback from customers.

Furthermore, a vendor’s transaction history reduces the likelihood of market exit as a longer transaction history is correlated with continued market participation. In licit online markets, the spectre of negative feedback also looms large. Cabral and Hortacsu (2010) demonstrate that when an eBay vendor receives their first negative feedback their weekly sales growth decreases from +5% to –8%. Relatedly, research on Yelp by Luca (2011) has shown that online restaurant reviews impact restaurant demand, especially for independent restaurants, a result which has been confirmed for hotels by Hollenbeck (2017). As well, Wagner (2016), in a field experiment in a Chilean start-up accelerator, found that negative feedback decreases the probability of start-up’s continuation, i.e. increases the probability of exiting the market.

Using longitudinal data from the first Silk Road, Przepiorka, Norbutas, and Corten (2017) examined the benefits of a good reputation. Moreover, the author examined how much buyers take into account sellers’ reputations when deciding whom to buy from. The authors found that vendors react to changes in their reputation by adjusting the prices of their goods. Vendors with a high reputation score would routinely increase the prices of their products as their devoted clientele would continue to return to them. The same cannot be said for less reputable vendors who decreased their prices to attract prospective buyers. This is phenomena was also documented by Shapiro (1983) in offline markets where low-rated sellers decreased their prices to compensate potential buyers for the risk they took when doing business with them. Interestingly, Przepiorka, Norbutas, and Corten (2017) also reported that higher rated vendors typically remained on the market for a longer period of time. This makes intuitive sense as a good reputation breathes longevity in one’s business as old customers will constantly return and new customers will join. Indeed, Bhaskar et al. (2017) maintains that online black markets manage to alleviate moral hazard problems predominantly because negative feedbacks lead to sales reductions. Importantly, vendor

reputation is transferable across markets. Norbutas, Ruitera, and Corten (2020), analysing vendor migration in three cryptomarkets, found that vendors that have accumulated a high cumulative reputation over many successful transactions were better able to migrate to another market following closure.

Importantly, this feedback system is not the only method by which vendors can earn the trust of buyers. Cryptomarket vendors will utilize a bevy of tactics to shore up their reputation. According to Christin (2013), vendors will, at the very minimum, address potential customers using a warm and professional tone. “The tone and content of these messages contrast significantly with the communication styles stereotypically associated with conventional drug dealers, and are likely to strike a reassuring chord amongst consumers who are accustomed to high levels of retail service in other aspects of their lives” (Martin, 2014a, 37). Vendors may also resort to licit retail techniques like Bitcoin lotteries and holiday sales. In fact, 4/20 or “International Pot Day” marked quite a celebratory affair on Silk Road as Ross Ulbricht waived all commission fees for marijuana purchases to reduce the overall consumer price (Martin, 2014a, 37). Furthermore, a study by Ladegaard (2017) revealed that new vendors seeking to cultivate a consumer base would offer low-cost and free samples.

Quite remarkably, the desire for a positive reputation is such that vendors will sometimes engage in corporate mimicry to an extent bordering on the outlandish. Certain vendors will employ conscientious market rhetoric, professing their commitment to selling products which use “organic” ingredients. In some cases, vendors may even attempt to sway consumers by proclaiming that their products had been purchased from poor agrarian farmers as opposed to violent drug dealers (Martin, 2014a, 39). Of course, terrestrial drug vendors have been found to screen potential customers, incentivizing long-term clients by offering credit or discounts (Chalmers & Bradford, 2013; Jacques, Allen, & Wright, 2014). Finally, risk-taking appears to be a proven method for establishing a good reputation. According to Decary-Hetu (2016), a willingness to ship overseas on the part of vendors was associated with higher reputation scores and greater profit. Shipping internationally is generally considered a perilous activity as it increases the risk of detection when drugs move across international borders (Volery, Mueller, and von Siemens, 2013).

Cryptomarkets, due in part to their semi-public nature, provide information on numerous transactions. However, because these exchanges take place without the benefit of face-to-face interactions, it is especially difficult for participants to gauge both the trustworthiness of others and the overall quality of products. The problem of uncertainty and information asymmetry has been extensively examined by Akerlof (1970). According to Akerlof (1970), the risk of market failure increases when buyers are unable to inspect products and differentiate them before purchase. As such, repeated experiences with low-quality sellers decreases buyers’ expectations and willingness to pay for high-quality products. What emerges then is a “lemon market” where consumers possess less valid or reliable information about the quality of the goods relative to vendors; this is information asymmetry. According to Herley and Florenio (2009), the uncertainty created by low-quality vendors imposes a tax on every transaction conducted in the market. That is, high-quality vendors stand to make less as the presence of low-quality vendors both discourages buyers from engaging in transactions and drives down the price of goods and services.

As with many types of real-world exchange situations, a clear way to establish trustworthiness is through transparency and the provision of accurate information. In a simulation study of the effects of positive reputation systems in a licit online market, Whitmeyer (2000) found that the effects of different types of positive reputation systems often depended to a large extent on the proportion of cooperators in the population. That is to

say, more cooperators in a market decreases uncertainty and information asymmetry by providing accurate information on the quality of vendors. In general, research (Cook and Emerson, 1978; Kollock 1999; Yamagishi and Matsuda 2003) into exchange relations in social networks demonstrates that uncertainty reduces the likelihood that an actor will form a relationship with an exchange partner given the potential for exploitation. This, however, is not the case when actors actively trust their prospective trade partners.

To date, only one study has examined trust networks in cryptomarkets. Duxbury and Haynie (2017) examined the local and global network structure of a transactional opioid network on the dark web. Using exponential random graph modelling, the authors demonstrate that the opioid network was highly localized, segmenting into subgroups where a small number of vendors accounted for a large number of transactions. As such, the authors concluded that “vendors’ trustworthiness is a better predictor of vendor selection than product diversity or affordability, with buyers choosing to conduct repeat transactions with trusted vendors” (Duxbury and Haynie, 2017, 23). This produces a unique network structure that is characterized by localized subgroups of comparable size. Building off of this study, Duxbury and Haynie (2018) also contend that social commerce networks on cryptomarkets are based on preferential attachment, where highly desirable vendors attract a large base of customers (Diekmann et al., 2014; Stephen and Toubia, 2009). Networks that form through preferential attachment generally exhibit a degree scaling property (Barabasi and Albert, 1999).

Duxbury and Haynie’s (2017, 2018) findings are corroborated by Decary-Hetu and Quessy-Dore (2017). Measuring the loyalty of repeat buyers over time, the authors find that, on average, buyers make 60% of their purchases from the same cryptomarket vendor. Nevertheless, the authors note that “while repeat buyers may want to remain loyal to a vendor they are often forced to purchase from other vendors when their main vendor is unable to supply them with the products they want” (Decary-Hetu and Quessy-Dore, 2017, 87). Though buyers generally purchase from the same vendor, this is not to suggest necessarily that all vendors operate on a single market. Using data collected on eight cryptomarkets, Broseus et. al (2016) examine market diversification among Canadian vendors. This analysis revealed that most of vendors (80%) focused their activities on only one market. Furthermore, their presence on several cryptomarkets at the same time decreases when the number of cryptomarkets increases.

As it relates to trust and uncertainty, however, buyers’ decision to repeatedly engage with a single cryptomarket vendor is indicative of Coase’s theory of the firm. According to Coase (1937), when market transactions are expensive or risky it makes sense to form relationships which ultimately culminates in a firm rather than purchase resources on a random basis. It makes little sense to transact with anonymous vendor when there is considerable uncertainty about the quality of products unless there is no alternative. As Casson (2001) suggests, “a firm may be defined as a specialized decision-making unit, whose function is to improve coordination by structuring information flow, and which is normally endowed with legal privileges, including indefinite life” (58). In the case of cryptomarkets, persistent uncertainty is a stimulus for the formation of transactional subgroups (Duxbury and Haynie, 2017; 2018). Indeed, the formation of subgroups is premised on trust between a small number of vendors and their respective customers. As buyers will not often do business with vendor(s) they have little experience with, they will consistently return to their primary vendor(s) in order to conduct further transactions. This naturally equates to a fairly diffuse trade network where buyers and individual vendors cluster in silos. Indeed, the formation of trust among cryptomarket users ultimately determines the structure and composition of the transactional network therein.

Law Enforcement Intervention and Network Disruption

The Digital Age has certainly simplified a bevy of once complex functions of daily life (e.g. communication, transportation, etc.). However, such enthusiasm is not necessarily shared by law enforcement as they must now contend with the emergence of new technology-oriented crimes. Of course, the ungovernability and unpredictability of the Internet has created lucrative opportunities for criminals looking for a quick score. Indeed, the oversaturation of online criminal markets means that criminal groups are not operating in isolation as market competition begins to ramp up. Indeed, a reduction in illicit market shares may lead to a reduction in profit among competitors. In short, competition breeds revenue. Nevertheless, the success of cryptomarkets is indicative of a digital revolution in crime. This section will examine the efforts made by law enforcement to combat and curtail these illicit online marketplaces.

Law enforcement agencies certainly recognize cryptomarkets as a credible threat. Consider DEA Special Agent Cromwell's characterization of darkmarket operators as greedy criminals, cowardly hiding behind encryption technology in order to peddle products which cause the deaths of 200,000 Americans on a yearly basis (Martin, 2014a, 2014). This is very similar to DEA Acting Administrator Michele Leonhart's statement on Mexico's La Familia Cartel: "this organization, the newest of Mexican cartels, is directly responsible for a vast majority of the methamphetamine pouring into our country across our Southwest Border, and has had a hand in fueling the cycle of violence that is wracking Mexico today". While it cannot be argued that the scope and influence of cryptomarkets rival that of Mexican drug cartels, it is apparent that law enforcement officials are taking these illicit entities seriously.

Nevertheless, this proactive focus on cryptomarkets is quite remarkable given law enforcement's long documented disinterest in cybercrime. "There is an omnipresent undercurrent of social stigma against those who fulfil less dangerous duties in law enforcement" (Goodman, 1997, 479). Given the lack of serious violence and associated difficulty of detecting and apprehending cybercriminals, cybercrime is not as stringently policed as terrestrial crimes. However, it stands to reason that the continued interventions against cryptomarkets are an exception to the rule. Indeed, law enforcement interventions have created an insalubrious environment for cryptomarkets as several popular firms have folded in a short period of time (Martin, 2014a, 65). However, while these initiatives have produced victories for law enforcement, costly defeats have also accrued.

While typically slow in counteracting the emergence of new cyberthreats, law enforcement organizations have made several attempts against cryptomarkets in the past decade. The first cryptomarket arrest occurred in 2013 in Western Australian. What is interesting about this case was that the offender was a local dealer seeking to resell purchased product on his own turf (Martin, 2014a, 58). To this extent, Aldridge and Decary-Hetu (2014) maintain that these "business to business" transactions represent approximately 31% to 45% of Silk Road sales revenue. This is further corroborated by Norbutas (2018) who notes that cryptomarket distribution networks are localized whereby vendors and buyers from the same countries typically do business with one another. Moreover, while there is no evidence to suggest widespread collusion between cryptomarket vendors and terrestrial drug traffickers and dealers, law enforcement agencies have certainly considered the influence of cryptomarkets in their nation's domestic drug supply.

As it pertains to law enforcement intervention in cryptomarkets, however, the scholarly literature has exclusively examined the effect of market takedowns and infrastructural disruptions. On October 2, 2013, the FBI-led arrest of Ross Ulbricht resulted

in the shutdown of Silk Road and the seizure of over USD \$33 million in bitcoins. Many Silk Road participants migrated to other markets following its closure. In fact, Soska and Christin (2015), analysing two years of transactional data, show that within six weeks of Silk Road's shutdown the number of dealers on Black Market Reloaded and Sheep increased by 200% and 400%, respectively. Moreover, by late January of 2014, sales volumes on several cryptomarkets exceeded what was documented on Silk Road. As such, this takedown cannot be described as being successful in the long term.

Bhaskar, Linacre, and Machin (2019), examining over 1.5 million drug sales, note that sales listings on Sheep market rose from 4,358 on October 17, 2013 to 8,457 by October 30, 2013. By April 2014, there were a combined 32,000 drug listings on Silk Road 2.0, Agora, and Evolution, 128% higher than the original Silk Road. The authors (2019) conclude that "there is no evidence that these exits deterred buyers or sellers from online drugs trading, as new platforms rapidly replaced those taken down, with the online market for drugs continuing to grow." Within two to three months of shutdown, vendor activity and consumer confidence returned to normal, with the overall market reverting to equilibrium. Furthermore, Buxton and Bingham (2015) found that, following the Silk Road shutdown, participants adopted more secure communication and encryption techniques. This particular finding is important as it suggests that the tactics and technologies used by cryptomarket users improves with each market closure law enforcement intervention. This equates to a game of brinkmanship where law enforcement must continually improve their capabilities in order to keep pace with cybercriminals on the dark web. These prospects do not necessarily bode well for law enforcement given constraints on resources.

Based on the available evidence, it can be argued that the closure of the Silk Road made policing a more difficult task as opposed to an easier one. This iatrogenic effect is readily observable in the myriad of cryptomarkets that emerged following the takedown of Silk Road. Martin (2014a) notes that "this is due partly to the fact that cryptomarket trading is significantly more decentralised now than it was when Silk Road was operating at its peak" (13). By late January 2014, sales volumes on several cryptomarkets exceeded what was documented on Silk Road. This explosion in sales and new markets has been hailed by many as irrefutable proof of the so-called 'hydra effect' (Ormsby, 2014). Indeed, the removal of one cryptomarket gives rise to many new ones.

Following the arrest of Ross Ulbricht, the second major disruption came in November 2014 when the FBI in collaboration with the Department of Homeland Security and Europol initiated Operation Onymous, shut down multiple cryptomarkets and arrested many users worldwide (Barratt and Aldridge, 2016). Examining the longitudinal impact of Operation Onymous, van Buskirk (2017) observed temporary decreases in vendors and listings, with the rate of vendor numbers increasing at constant rate. However, van Buskirk maintains that "as of November 2015, the overall number of vendors had not returned to the level seen in November 2014, just prior to Operation Onymous". This finding is challenged by Decary-Hetu and Giommoni (2017) who measured supply side indicators across five cryptomarkets in the 41 weeks that preceded Operation Onymous and the 21 weeks that followed it. The authors found that initial decreases in market activity were entirely offset by long term gains. That is to say, while the number of listings and vendors decreased in the first several weeks following Operation Onymous, they recovered entirely in the following months. In fact, sales doubled as early as two months following the intervention. As Bhaskar, Linacre, and Machin (2019) maintain, "overall, it is not possible to find evidence of deterrent effects associated with either the two law enforcement shutdowns" (230).

Comparable to operations against terrestrial drug trafficking organizations, the efforts made by law enforcement in curtailing cryptomarkets have yielded less than desired results. The criminological literature (Kenney 2007; Gambetta 2009; Malm and Bichler, 2011) has documented the divarication of organized crime entities following successful state intervention. To this extent, the absence of a monopolistic entity creates a vacuum of unappeased demand for which smaller entities will scramble to fill. On a national level, the Columbian government's assault on the Medellin and Cali cartels had thoroughly fragmented the drug market to a nearly unmanageable degree. "Following the DEA's kingpin strategy in the 1990s, during which US and Colombian law enforcers effectively decapitated the most notorious 'cartels', numerous so-called 'micro-cartels' emerged in their place" (Kenney, 2007, 257). On a local level, the NYPD's arrest of Nicky Barnes in the 1970's had the inadvertent effect of fragmenting the New York heroin market. "In destroying the Barnes monopoly, law enforcement practices created . . . an opening in the market that was filled by new distributors, who literally wanted to make a name for themselves in order to increase their share in a burgeoning market" (Gambetta, 2009, 202). Of course, this "hydraization" is also present on the dark web.

In 2017, coordinated law enforcement operations saw the closure of two large drug cryptomarkets: Alphabay and Hansa. However, according to Afilipoaie and Shortis (2018), the strategies used in this operation differed from previous interventions as they were intended to damage the trust which undergirds business-to-consumer relations rather than simply close the marketplace. To elaborate, the FBI closed AlphaBay without posting a seizure notice or making a public statement so as to allow users to flock to Hansa, which saw an eight-fold increase in users. However, Hansa had been co-opted and secretly ran by the Dutch National Police prior to AlphaBay's closure. U.S and Dutch official, together with international partners, then initiated a "knock-and-talk" operation on addresses they had secured from the bust (Aldridge, and Barratt, 2020). Users were visited at their homes and warned against using cryptomarkets in future. In some cases, arrests were made. It is, however, unclear what impact this intervention has had.

Early research demonstrates that whilst users from Alphabay migrated to Dream Market in a similar pattern to previous takedowns, users from Hansa opted instead to change their PGP keys or usernames, suggesting they chose security over maintaining their marketplace reputations (Van Wegberg and Verburch, 2018). However, findings from the Internet Institute indicate that the overall cryptomarket trade volume returned to pre-bust levels within a month of Alphabay's closure (Dittus, 2017). While this intervention has not produced the results perhaps desired by law enforcement, its sophistication relative to earlier operations should be viewed as a positive outcome. With each strategy, law enforcement are perhaps becoming increasingly knowledgeable as to what works and what does not, adapting and adjusting the parameters of future strategies to incorporate lessons from prior interventions. However, this is contingent on whether law enforcement dealing with cryptomarket are made aware of the measurable impact of their interventions. This raises questions about the use an evidence-based calculus when policing cryptomarkets. Are the results of past interventions used to determine how future interventions are structured? Based on the available evidence, this may not be the case.

It appears that law enforcement interventions against cryptomarkets have been ineffective and perhaps counterproductive. In the aftermath of market closure, sales volumes generally returned to comparable pre-closure levels while new markets emerged to take the

place of those that were shut down. This is the general pattern. It is important to note that this fragmentation is partially due to decentralized exchange networks of cryptomarkets. “In the event that a cryptomarket is shut down, the user community is able to persist; users either migrate to other sites or, as in the case of Silk Road 1.0, they construct and quickly repopulate a replacement website” (Martin, 2014a, 23). This mobility and durability equates to a difficult-to-exterminate illicit entity. Though it is perhaps reprehensible to allow the unabated operation of organized crime, it is arguably far worse for law enforcement to destroy a criminal monopoly as the crime problem is allowed to metastasize at a greater rate.

It is reasonable to conclude that these interventions have had an iatrogenic effect, facilitating the growth of cryptomarket activities to levels greater than pre-intervention operation. In short, law enforcement interventions against cryptomarkets have produced short term gains, temporarily disrupting the ease for operation of these illicit platforms and deterring vendors and buyers from continued operation. However, in the long term, these law enforcement interventions have paradoxically made policing the dark web a more difficult task as more cryptomarkets with greater risk reduction competencies have emerged. These markets have grown larger, generating more revenue while catering to an increasing number of vendors and buyers. Moreover, specialized markets which cater to specific customers have both emerged at a greater frequency and have gone further underground, away from the prying eyes of law enforcement monitoring the dark web.

What is particularly telling is that Silk Road has itself undergone several resurrections following closures, returning as Silk Road 2.0, Silk Road 3 Reloaded, and the latest iteration Silk Road Reloaded. One can only imagine as to what the current dark web environment might be like had the Silk Road’s monopoly been kept intact. Indeed, it is not outside the realm of possibility that while new cryptomarkets might have emerged to compete with Silk Road, the sophistication and profit-maximization of these platforms might have been far lower than they are today.

While it is imminently clear that largescale market closures are not the way forward, there is an open question as to what is. Scholars have increasingly focused on the network dynamics within cryptomarket transactional networks as a means of understanding their structural vulnerabilities. As with studies in this particular subfield, these studies seek to identify the structural vulnerabilities in a cryptomarket as well as the strategies which might best take advantage of these vulnerabilities. This is a potentially fruitful avenue of research as the results might serve to inform strategic decision-making when it comes to cryptomarket interventions. Duxbury and Haynie (2018; 2019) have made the most progress in area, applying adaptive computer simulations to test the theoretical effect of law enforcement interventions on a cryptomarket transactional network.

Building off their prior work (Duxbury and Haynie, 2017) on the network structure of a cryptomarket, Duxbury and Haynie (2018) conducted disruption simulations on the same opioid market. In particular, the author’s identified vendor selection patterns using exponential random graph models then evaluated the network’s robustness using vertex removal simulations. Given that this opioid network was characterized by degree scaling properties pursuant to preferential selection of vendors on the part of buyers, the size and scope of the market was reduced with the sequential removal of the top vendors therein. To this extent, the size of the largest components shrank while the proportion of potential

components and number of isolates in the network decreased and increased, respectively, as more vendors were removed. This study demonstrates two interrelated principles with regard to the network structure of cryptomarkets. First, Duxbury and Haynie (2018) observe that the evidence of preferential attachment mechanisms “lends greater support to the influence of trust than the effect of product differentiation or affordability” (246). Second, this trust can be exploited by interventions seeking to disrupt a cryptomarket’s ease of operation.

In their second study, Duxbury and Haynie (2019) designed an agent-based simulation to assess the network responsiveness of a larger darknet drug market. The authors considered three attack strategies: 1) weak link attacks that delete large numbers of weakly connected vertices, 2) signal attacks that saturate the network with noisy signals, and 3) targeted attacks that delete structurally integral vertices. The authors demonstrated that targeted attacks generally succeeded in disrupting the market when adopted at a large scale. The authors (Duxbury and Haynie, 2019) conclude that that “these two processes undermine long-term network robustness and increase network vulnerability to future attacks”.

It is important to emphasize that these results should not be accepted dogmatically given the nature of adaptive computer simulations. Scholars leveraging adaptive computer simulations are merely making educational guesses on the assumed rational behaviour of actors in a criminal network. As such, modelling parameters are based on these assumptions. Whether cryptomarket actors behave in this manner is another matter altogether. In short, while adaptive computer simulations go some way towards identifying structural vulnerabilities in cryptomarkets, they should not be accepted as the complete truth. The behaviour of licit actors much less criminal actors cannot be perfectly simulated given the probabilistic nature of human behaviour. While general patterns in human behaviour are observable, strict obedience to these patterns will differ from actor to actor.

Regardless, the results of these studies are promising for designing effective law enforcement strategies to combat cryptomarkets. Adaptive rule-based sequential node removal goes some way towards mimicking the operation of a cryptomarket when pressed by a targeted intervention. Law enforcement might find use in applying this methodological technique when deciding which actors to taken and how the removal of these actors might affect the overall structure and operation of the market. However, there is a pressing need for more studies which simulate law enforcement interventions on real-world cryptomarket transactional networks in order to evaluate the impact of specific targeting strategies. In particular, such studies should test the efficacy of individual targeting strategies, determining their ability to disrupt the operation of a cryptomarket and how this performance stacks up against other targeting strategies. Furthermore, these studies must incorporate some form of network adaption to mimic the purported behaviour of actors when the market is disrupted. Given that criminal networks are comprised of human actors whose behaviour is liable to change in the face of an attack, studies which leverage computer simulations to understand the impact of strategic interventions must consider probable adaptation on the part of actors within the network.

Conclusion

This chapter has attempted to consolidate the cryptomarket literature, identifying all primary strands of the research to-date. Of course, the relative novelty of cryptomarkets means that the scholarly literature on cryptomarkets is still in its infancy. There is, indeed,

more work which needs to be done. However, if one were to describe the historical transformation of cryptomarket research, it would be useful to segment the literature into three distinct phases:

- 1) Products, places, and people
- 2) Trust and network structure
- 3) Network robustness and strategic interventions

The first phase of cryptoamarket research can be construed as an exploratory phase where early cryptomarket researchers sought to understand what products were sold on these illicit entities, who bought and sold them, and which countries they were shipped to and from. As such, this particular phase of the research attempted to document the basic or perhaps superficial elements within cryptomarkets. There was also a greater emphasize on the use of descriptive statistics and qualitative methods to understand what these entities were and how they operated. The second phase of cryptomarket research pertains to studies examining the formation of trust on cryptomarkets and the untangling of the network structure of these illicit entities. To this extent, more sophisticated research methods, including social network analysis and statistical modelling, were used to determine the processes by which actors on cryptomarkets came to trust one another and how this trust is carried over into future transactions. There was, moreover, an explicit focus on the rank and position of cryptomarket pursuant to their reputation. Finally, the third and current phase of cryptomarket research concerns the examination of the robustness of cryptomarket transactional networks as well as evaluations of the law enforcement interventions. This research leverages computer simulation methods to answer questions about the efficacy of strategic interventions against cryptomarkets. Furthermore, these analyses attempted to clarify which tactics and strategies worked and which ones were less than successful.

It is important to stipulate that these phases are not mutually exclusive as descriptive research on cryptomarkets still persists today. Moreover, earlier phases of the research have not ceased as there are still a number of research questions which must still be answered. For example, the proliferation of synthetic opioids has generated studies on the use of cryptomarkets in trafficking these substances. What is evident from this historical transformation is the ever-increasing level of methodological sophistication featured in studies. This is perhaps reflective of the depth and quality of the research questions being asked. The proposition of ever-ambitious research questions and objectives requires the use of increasingly sophisticated research methods. It is, however, an open question as to where the research will go in the coming years. Moreover, there are both slight and substantial gaps in the scholarly literature which must be filled.

In this regard, there are several pressing questions which must be asked and topics which must be examined by cryptomarket scholars in order for the literature to progress. While data which links buyer and vendors together via unique identifiers is relatively scarce, there is still a pressing need for research which examines the network structure and robustness of cryptomarkets. Indeed, studies by Duxbury and Haynie (2017) and Norbutas (2018) must be replicated for us to determine generalizability of their findings. Furthermore, more studies which measure the structural robustness of cryptomarket transactional networks and the associated efficacy of strategic interventions are required. Indeed, the research has documented the operational elements of these entities but must now veer into more practical matters. This involves determining the strategies and tactics which might best disrupt the ease of operation of these illicit entities.

To this extent, cryptomarket research is bereft of experimentation of any kind. This pertains to testing and tracking the efficacy of interventions in real time via control and treatment groups. While the logistics of such research is understandably complex, it is a necessary step forward in the domain of evidence-based cryptomarket research. Current research examining strategic interventions rely on adaptive computer simulations. While this is certainly useful in matters of theorization, carefully designed experiments are a step up, providing actionable intelligence on the effectiveness of strategic interventions against cryptomarkets.

Furthermore, while administrators such as Ross Ulbricht have been arrested, it is unclear who exactly establishes cryptomarkets and, more importantly, what their motivations are. Current research has examined the demographics and motivation of buyers but has yet to do the same for those who operate these illicit entities. Moreover, it is unclear how cryptomarket administrators recruit moderators. As well, researchers have yet to examine in detail how cryptomarkets have innovated in response to law enforcement interventions, how fast these adaptations were made, and how effective they have been. This particular set of question deals with the innovative nature of cryptomarkets, an area which may aid law enforcement in understanding the potential outcomes of future interventions. Furthermore, it is unclear what role, if any, cryptomarkets play in the proliferation of new illicit drug trends. This is related to the increasing use of fentanyl and other dangerous synthetic opioids.

Cryptomarkets represent a fascinating area of study for researchers interested in the intersection of cybercrime and network science. In the following chapters, I endeavour to examine three topics in greater detail: the network structure and formation of trust on cryptomarkets, consumer satisfaction and information asymmetry, and the efficacy of targeted interventions.

Chapter 2: Trust under Uncertainty: How Network Structure and Vendor Selection Inform Trust Formation on Cryptomarkets

While advances in digital communication have yielded unprecedented opportunities for commerce and social engagement, it has also created new opportunities for crime and deviance. Indeed, cybercrime is one such area where stable increases in the complexity and sophistication of crime is readily observable. Moreover, given the immaterial nature of computer-enabled offenses, those looking to collaborate need not gather in a physical location. Rather, prospective cybercriminals can collaborate from the comfort of their own homes, jointly hacking government websites (Lusthaus, 2018) or discussing the latest techniques for committing offenses without being detected much less apprehended (Decary-Hetu and Dupont, 2013). In short, cybercriminals are taking advantage of technological advancements for the purpose of collaborating in committing crimes (Lusthaus, 2018).

This is particularly the case for illicit online marketplaces hosted on the dark web. These cryptomarkets, as they are called, function as brokerage platforms, connecting capable vendors and willing buyers looking to truck, barter, and trade in a variety of illicit goods and services. Owing to their relative success and continued growth, these platforms mimic the structure, operation, and financial risk competencies of licit platforms such as eBay and Amazon. In other words, they provide the necessary structure and order that is often missing in terrestrial criminal markets. Nevertheless, the communal nature of cryptomarkets raises questions about the behaviour of the actors therein. Indeed, there are open questions about how buyers identify and select vendors, how these transactional relationships change over time, and how this ultimately affects the network structure of the market. Understanding the transactional network of cryptomarkets is necessary if we are to answer these questions. Moreover, intelligence on the network structure of a cryptomarket may provide crucial insight into the vulnerabilities therein. This possesses practical implications for law enforcement organizations attempting to disrupt the ease of operation of these illicit entities (Bright et al., 2017).

Following work done by Duxbury and Haynie (2017) and Norbutas (2018), this study examines the network structure of Abraxas, a cryptomarket in operation between 2014 and 2015. It will, furthermore, identify the market-level metrics that predict for vendor selection as well as the developmental trajectory of vendor performance. Together, these results provide further insight into how trust among buyers and vendors determines the structure of cryptomarkets.

Following Papachristos (2009; 2014) and Duxbury and Haynie (2017), I employ social network analysis to both construct and analyse this transactional network. Over the past two decades, an increasing number of studies have leveraged social network analysis (SNA) to understand the inner workings of various covert networks (Holt, Strumsky, Smirnova, & Kilger, 2012; Kenney, 2007; Morselli, 2009; Malm & Bichler, 2011; Natarajan, 2006; Wood, 2017). I also apply community detection analysis to determine the underlying subgroup structure of this cryptomarket. Finally, I employ statistical modelling and trajectory modelling to determine which factors which predict for vendor trustworthiness and the developmental trajectory of trusted vendors. As such, this study seeks to replicate social network analyses conducted by Duxbury and Haynie (2017) and Norbutas (2018) while

offering novel contributions relating to the predictors and developmental trajectory of vendor trustworthiness.

This combination of descriptive network analysis, community detection analysis, and statistical and trajectory modelling allows for a thorough examination of trust formation and network structure on the Abraxas cryptomarket. Fundamentally, this study seeks to test the generalizability of findings made by Duxbury and Haynie (2017) but seeks also to improve upon this work by examining additional explanatory factors and the longitudinal performance of vendors. In general, this research seeks to disentangle trust dynamics on a dark web market, uncovering the processes by which trust is created and maintained and how this ultimately affects the network structure of the market.

Literature Review

Trust in the Criminal World

Within criminal enterprises and associations, trust is a fundamental but difficult-to-establish operational tool. Trust, in other words, is a fragile component within criminal undertakings involving more than one actor. This is primarily due to the uncertainty associated with anti-social and deviant behaviour. Indeed, the situational constraints with which a criminal must contend (death, arrest, betrayal, etc.) often encourages these actors to renege on stated or perceived obligations. Moreover, there is generally no principal authority which can uphold contractual obligations and punish dissenters in illicit environments as would be done in licit settings. Though organized crime groups such as the Italian mafia (Gambetta, 2000; Catino, 2014; von Lampe, 2016) often engage in some form of governance over the entities it presides over, this is a relative rarity in the criminal world. Still, trust is a coordination tool, allowing criminals to cooperate and strive toward a common objective.

How is trust defined within a crime context? Many scholars have offered a definition. Gambetta (2000) defines trust as “a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action, both before he can monitor such action (or independently of his capacity ever to be able to monitor it) and in a context in which it affects his own action” (217). Von Lampe and Johansen (2004) note that trust is a mechanism for individuals to “cope with risk and uncertainty in interactions with others” (103). Dumouchel (2005) defines trust as “an expectation concerning another agent's action that is relevant to the decision to act” (421). For these scholars, trust involves the presupposition of future risk. It, thus, requires an actor to ascertain, to the best of their ability, the interests and predilections of the those whom they intend to engage with. This involves determining how a trustee might behave in a particular situation. To this extent, Gambetta (1988) notes that to bestow one's trust in another actor requires an active consideration of the subjective probability that this agent may betray you or fail to uphold their part of the contract in some capacity.

However, trust is difficult to establish in the criminal world as it requires the suspension of selfish desires on the part of self-interested actors (Williamson, 1993). This amounts to a trust deficit in the criminal world; an outcome which most affects those seeking to place trust in others. As Gambetta and Bacharach (2001) demonstrate through a game theoretic approach, the optimal outcome for a trustee is to cheat (renege) in the event that a

truster opts to cooperate (endow trust). However, Gambetta (2009) also contends that the iteration of this outcome over several rounds would be most detrimental to a trustee. That is, their persistent duplicitousness would discourage the truster from cooperating, costing the trustee all future business opportunities (38). In this regard, trust involves knowing whether those whom you engage with are sincere in their intention to cooperate or are merely feigning their cooperation and instead playing an altogether different game. This trust deficit is further compounded by the fact that long-term criminal partnerships require a high level of trust. As such, if trust cannot be consistently maintained, these partnerships will be abrupt and sporadic.

Still, it is important to stipulate that cooperation in the criminal world can occur without trust. In events typified by a negative-sum outcome (both actors stand to lose), the establishment of trust is not required. In these situations, agents will act out of mutual self-interest as failure to do so may likely result in sanctions against all agents involved. As such, an agent need not explicitly bestow their trust in another agent much less engage in presuppositions of future risk as it is self-evident that the opposing agent is acting out of interests which coheres with one's own interests. However, such a schema is built upon two requisite elements: 1) all agents are aware of their own interests and 2) all agents can confirm that their interests align with the interests of other agents. While the first element is perhaps easy to establish as it requires knowing which outcomes one desires, the second element may prove problematic to establish as an agent cannot always discern the outcomes which are desired by other agents. However, this is not to suggest that trust is absent in these particular situations as it manifests in a different form. While an agent may not trust a prospective partner, he or she can trust their intentions.

While the criminological literature (Gambetta, 2000; von Lampe and Johansen, 2004; Gambetta, 2009; Campana and Varese, 2013) has emphasized the trust deficit within criminal networks of varying size, these observations reflect criminal activities which take place in terrestrial markets. One would be correct in assuming that trust dynamics are liable to change in cyberspace. Several studies (Holt and Lampke 2010, Yip 2011) have indicated that market-driven dynamics are present in illicit online markets. That is, illicit online exchanges are treated more like voluntary economic transactions than they are mere illicit transactions. Decary-Hetu and Dupont (2013), in an examination of a botnet forum, found that simple indicators like the number of awards received, number of days spent on the forum, or the size of one's network often determined how well a vendor was trusted. In this case, surface-level trust was built upon personal characteristics and behaviour as opposed to mutual experiences where a deep trust could be developed.

Trust and Reputation on Cryptomarkets

Relative to terrestrial markets, cryptomarkets are bastions of collegiality and cooperation. While this is not to suggest that duplicity and deception are absent from these platforms, cryptomarket participants are generally more trusting of one another than are participants on illicit terrestrial markets. Van Hout and Bingham maintain that the relationships of cryptomarket participants were "based on levels of trust and professionalism" (Van Hout & Bingham, 2013, 387). This is due primarily to the manner in which information is shared on cryptomarkets. To this extent, vendors openly share information on the quality of the goods and services they sell whereas buyers provide publicly accessible feedback on their experience with these vendors. As such, the quality of a good or service and the

trustworthiness of a vendor can be more easily discerned on cryptomarkets than in offline markets. In fact, a study conducted by Dasgupta et al. (2013) found that buyers were able to “provide a valid estimate of the street price of diverted prescription opioids...and predict the relative pharmacologic potency of opioid molecules” (178).

On cryptomarkets, vendor reputations are created by repeated transactions with buyers who rate their experience with a specific vendor. This is based primarily on a numerical feedback score (e.g. 0 to 5 stars) but also includes written feedback which offers greater detail on the transaction. A cryptomarket vendor cannot alter the feedback published on their page, whether positive or negative (Martin, 2014a, 42). As such, reputations cannot be artificially inflated by self-serving vendors as they are organically created by transactions with buyers. Tzanetakis, Kamphausen, Wersé, and von Laufenberg (2016) argue that customer feedback in cryptomarkets creates trust in an environment which is often bereft of it. Indeed, the illegal drug trade is often without assurances of the actions and intentions of one’s prospective trading partner(s). To this extent, a vendor’s reputation constitutes public knowledge on these platforms as prospective buyers can access it by simply visiting a vendor’s page and reading the vendor’s overall reputation score as well as the comments left by past buyers.

Hardy and Norgaard (2016) use data on cannabis listings from Silk Road to study the relationship between reputation and prices. The authors show that “reputation acts as a sufficient self-enforcement mechanism that allow transactions to occur” (Hardy and Norgaard, 2016, 32). To this extent, vendor reputations constitute a formal institution that creates a stable trading environment among those least expected to deal honestly. Janetos and Tilly (2017) show that a mature highly-rated cryptomarket vendor charges 20% higher price than a mature low-rated vendor. In general, vendors with more reviews charge a higher price than sellers with a low number of reviews regardless of rating. However, bad (i.e. low-ranked) sellers prefer to exit the market than decrease their prices in response to negative feedback. This is in line with Batikasa and Kretschmera (2018) who, studying the Agora marketplace, found that cryptomarket vendors are more likely to exit following negative feedback.

Duxbury and Haynie (2017) examined the local and global network structure of a transactional opioid network on the dark web. They found that the cryptomarket transactional network was diffuse and highly localized, with many buyers doing business with a small number of vendors. As such, the transactional network consisted of numerous subgroups based around several popular and prosperous vendors. These localized subgroups were of comparable size. With regard to trust, the authors concluded that “vendors’ trustworthiness is a better predictor of vendor selection than product diversity or affordability, with buyers choosing to conduct repeat transactions with trusted vendors” (Duxbury and Haynie, 2017, 23). Building off of this study, Duxbury and Haynie (2018) also contend that social commerce networks on cryptomarkets are based on preferential attachment, where highly desirable vendors attract a large base of customers (Diekmann et al., 2014; Stephen and Toubia, 2009). The structure of these networks is premised on a degree scaling property where a small number of nodes share ties with many other nodes within the network (Barabasi and Albert, 1999). As such, based on these findings, cryptomarket transactional networks are governed by a degree scaling property.

Norbutas (2018), examining the transactional network of the Abraxas cryptomarket, made similar findings to Duxbury and Haynie (2017). In particular, the author found that the

Abraxas transactional network exhibited low network density, with a small number of vendors accounting for the majority of transactions. Using exponential random graph modelling, Norbutas (2018) also demonstrated that Abraxas' transactional network was highly localized, segmenting based on geographical considerations. As such, the author concluded that the structure of Abraxas' transactional network was governed by geographical boundaries where vendors generally shipped to buyers from the same country. This is in contravention of the popular belief that cryptomarkets are multi-national entities where transactions occur between actors from different parts of the world. In contrast, cryptomarkets may instead solidify domestic trading, keeping illicit products within the borders of a nation.

Research Questions

This paper seeks to answer four research questions:

1. What is the network structure of Abraxas?
2. What is the composition of transactional communities within the network?
3. What market-level metrics and/or vendor characteristics predict for vendor trustworthiness (i.e. success (completed transactions), popularity (unique buyers), and affluence (revenue))?
4. What is the developmental trajectory of vendors' success, popularity, and affluence during their tenure on Abraxas?

Given our understanding of the value of trust in cryptomarkets, it is perhaps natural to ponder about the contexts in which transactions occur. Indeed, we have yet to fully understand the network structure of cryptomarkets and how this may be associated with decision-making processes under conditions of uncertainty. Furthermore, what is not well understood are the variables which predict for the selection of cryptomarket vendors by buyers. Indeed, the issue of preferential selection among buyers is especially curious in light of the problem of information asymmetry (Akerlof, 1970). According to Akerlof (1970), the risk of market failure increases when buyers are unable to inspect products before purchase. What emerges then is a "lemon market" where consumers' lack of valid or reliable information about the quality of a good or service sold imposes a tax on every transaction conducted on the market (Herley and Florencio, 2009).

The first research question seeks to determine the global structure of Abraxas' transactional network. While Duxbury and Haynie (2017) have examined this particular phenomenon in another cryptomarket, dark web researchers are generally uncertain about how vendors and buyers orient themselves within the transactional network they inhabit. Certainly, it is the case that vendor reputations serve to distinguish high-quality vendors from low-quality vendors, but there may other unexamined factors. Moreover, it is unclear how trust affects the overall network structure of a cryptomarket. As Barratt and Aldridge (2016) highlight, research into the network structure of cryptomarkets can provide insight into hidden transactional dynamics that stabilize these illicit online marketplaces. As such, this first question seeks to build off of Duxbury and Haynie's (2017) research, examining the network structure of a second cryptomarket. In short, I seek to test the generalizability of Duxbury and Haynie's (2017) findings by applying their methods to another cryptomarket transactional network. It is also important to state that Norbutas (2018) has done similar work on the Abraxas cryptomarket. However, this study will offer a more in depth look at the structure of Abraxas' global transactional network.

A logical follow-up to the first research question, the second research question seeks to understand the characteristics and composition of identifiable communities within Abraxas. Analyses by Duxbury and Haynie (2017) demonstrate that cryptomarket users orient themselves into subgroups whereby single vendors transact with many buyers. As such, the cryptomarket transactional network is reminiscent of small islands that are product and country specific. Importantly, no other study has applied community detection to a cryptomarket transactional network. As such, more research is required on this particular area. Community detection analysis will aid in further understanding the network topology of cryptomarkets. As with the first research question, this question seeks to test the generalizability of Duxbury and Haynie's (2017) findings on another market.

The third research question seeks to identify the characteristics that best predict for vendor selection. While the current research (Decary-Hetu and Quessy-Dore, 2016) can tell us which vendors are popular, it has yet to tell us why this is the case. Understanding how buyers select vendors is critical for understanding how the network structure of a cryptomarket comes to be. This question deals primarily with trust. In this case, I am attempting to quantify which market-level metrics predict for vendor selection across three proxy variables for trust. Following Gambetta (2000), trust, for the purposes of this paper, is defined as "a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action". As such, the proposed market-level metrics may serve as indicators (or game theoretical tools) by which buyers assess whether or not a vendor will uphold their end of an established transactional agreement. This study contributes to the literature by measuring 14 predictors across three different conceptualizations of vendor trustworthiness. This qualifies as the most extensive undertaking to date.

This fourth research question seeks to understand the developmental trajectory of cryptomarket vendors, assessing whether the most trustworthy vendors remain prosperous as the market expands. In this regard, what is not well understood among cryptomarket scholars is the extent to which vendors operating on these platforms remain at their current station and how they might grow or decline as they continue to operate on a market. This question offers insight into how market growth is affected by vendor growth and vice versa. If it is the case that a small number of trusted vendors are responsible for the majority of activity on a cryptomarket, it stands to reason that the continued operation and growth of this market is contingent on the performance of a core group of vendors. The bequeathment of trust upon vendors by buyers is thus a fundamental element by which a market is permitted to exist. This will serve as an entirely new contribution to the scholarly literature on cryptomarkets.

Data

Here I use a buyer-seller dataset from the Abraxas cryptomarket (Branwen et al., 2015). Apart from the anonymous cryptomarket analysed by Duxbury and Haynie (2017; 2019), this is the only marketplace where unique identifiers are available for buyers. Importantly, Abraxas was previously used by Norbutas (2018) in an examination of the geographical distribution of transactions. For my purposes, I construct a two-mode buyer-seller trade network with information on 5434 trades of illicit goods and services between 269 sellers and 2794 buyers, over a period of 7 months in 2014–2015.

As Norbutas (2018, 93) indicates, the dataset collected by independent researcher Gwern Branwen (Branwen et al., 2015) contains information on multiple cryptomarkets and is known to suffer from incompleteness. To specify, the entirety of the Abraxas marketplace might not have been captured in daily scrapes conducted by Branwen. To this extent, Norbutas (2018) compared the number of crawled item pages in these data to the actual number of items displayed in the home page of Abraxas at each date and found clear inconsistencies. More generally, Norbutas reported that “the average percentage of collected items across all of Branwen’s crawls is 92.4%, ranging from 26% to 100% depending on the crawl” (2018, 93). Furthermore, many of the scraped webpages are broken, meaning that the full extent of market transactions could not be recorded. This is a clear limitation as only a portion of the Abraxas cryptomarket could be examined. To this extent, this is not a complete transactional network as all transactions were not scrapped or re-recorded. Following Norbutas (2018), I aggregated information across all daily crawls of item pages. As a result, duplicate transactions were identified and removed. The aggregated data contains 269 unique sellers, 2794 unique buyers, and 5434 total transactions.

In order to construct a two-mode transactional network of exchanges between individual buyers and sellers, each collected feedback message needed to be attributed to a particular buyer. In general, feedback serves as documentable proof that a transaction has occurred. “Customer feedback takes a variety of forms, ranging from detailed comments about shipping times, ‘stealth’ measures and the perceived potency of illicit drugs, to a simple 5-star rating” (Martin, 2014a, 41). Importantly, while all cryptomarkets are feedback-based, they may differ on policies regarding the mandatory nature of buyer feedback. That is to say, some cryptomarkets require buyers to leave feedback after every transaction while others do not. Abraxas falls into the former category, with all transactions conducted over the market’s operational period being documented via buyer feedback.

While feedback data would ordinarily pose a problem in many network-based cryptomarket datasets due to partial or completely anonymized buyer nicknames, Abraxas contained unique buyer profile identifiers for each feedback message, which was located in the HTML code of item pages. I used these buyer identifiers to aggregate feedback messages left by each buyer account. Following the removal of duplicates, this permitted for the creation of a two-mode transactional network for vendors and buyers operating on Abraxas between January 15th, 2015 and July 4th, 2015.

While I was able to identify purchases made by individual buyer accounts, the data did not include buyers’ country of residence. Although I cannot observe buyers’ geographic location directly, inferences about geographic clustering in the marketplace can be drawn based on buyers’ selection of sellers located in particular countries. All transactions were organized into a variety of categories for analysis. These include a general category for all item types, a subcategory which disaggregated the items into more precise categories, and a secondary sub-category which provided more granular information on each item. Each item was hand coded. As it relates to pricing, all transactions were converted from bitcoin to USD based on a moving U.S. exchange rate. While this method might, in theory, produce less accurate pricing data given the volatility of cryptocurrencies, the listing prices also change as a result. As such, setting a fixed exchange rate, as opposed to a moving one, would not properly capture changes in listing prices.

Methods

Given its simplicity, transparency, and accessibility, descriptive statistics were used to summarize the 5434 transactions. This is done to understand both the nature and composition of illicit transactions on Abraxas. In general, descriptive statistics provide a clear and concise summary of the data. Importantly, social network analysis was also conducted to examine the network structure of Abraxas. In particular, I employ four analytic strategies: descriptive network analysis, community detection analysis, statistical modelling, and trajectory modelling. All network statistics, modelling, and visualizations were conducted in R and Microsoft Excel.

Descriptive Network Analysis

Given the use of social network analysis, standard network measures will be used to summarize the network structure of Abraxas at a cursory level. Importantly, I establish the presence of tie between two actors based on whether or not a feedback has been left from a transaction. The presence of feedback is documentable evidence that a transaction has occurred. Bichler, Malm, and Cooper (2017) correctly assert that researchers must clearly explain how they generated the networks for social network analysis. From the 5434 illicit transactions, a two-mode network featuring vendors and buyers was created. Only transactions with both a known vendor and buyer were used to construct this network. Vendors were identified based on their unique vendor name while buyers were identified based on their HTML code. As such, the transactional network consisted of 5434 transactions between 269 unique vendors and 2794 unique buyers. A link exists between actors if they were involved in a transaction together (McGloin and Kirk, 2011).

Here, I use four network measures: network density, in and out-degree centralization, and eccentricity. Density measures the interconnectedness of a network. To elaborate, this measurement divides the total number of ties between actors by the total number of ties which might be possible. There measurement is reflected by a coefficient which ranges between 0 and 1. As it relates to this data, a score close to 1 indicates that buyers do business with many vendors given the interconnectedness of the network. In contrast, density scores closer to 0 indicate that buyers transact with a small number of vendors and that the network is diffuse.

“Centralization measures how much influence a few actors exert over the network structure” (Duxbury and Haynie, 2017, 23). As it relates to this study, centralization tells us how vendors (outdegree centralization) or buyers (indegree centralization) influence network structure of the Abraxas transactional network. Centralization is determined by calculating the degree centrality of each node. “The sum of the differences between the actor with the highest centrality score and all other actors in the network is then divided by the largest possible sum of differences retrieved from a theoretical matrix of the same size” (Duxbury and Haynie, 2018, 929). This results in a value ranging from 0 to 1, with higher values indicating a greater central tendency in a network (Wasserman and Faust 1994). Finally, eccentricity measures the maximum distance of one node to any other node in the network. As such, the eccentricity of a node in a connected network is the maximum distance between that specific node and all other nodes in the network.

Each of these measurements were selected as means of determining both interconnectedness of the global network structure of Abraxas as well as the importance of individual nodes within the network. Other measurements such as closeness and betweenness centrality could have been applied but these measurements not have proved as insightful given the strict classification of each node as either a buyer or vendor.

Community Detection Analysis

While standard network measures provide insight into the aggregate features of a network, they do little in the way of unearthing underlying structural features within a network. This, however, can be achieved through community detection analysis. “Community detection refers to the procedure of identifying groups of interacting vertices (i.e., nodes) in a network depending upon their structural properties” (Yang et al., 2013, 15). In short, community detection algorithms will parse nodes into distinct communities based on the number of ties they have with other nodes within the network. Though exceptions exist, networks generally consist of actors who engage more regularly with some actors than they do others.

Here, I employ the walktrap community detection algorithm (Pons and Latapy, 2005; Newman, 2003; 2006) to determine the subgroup structure of the Abraxas transactional network. As Pons and Latapy (2005) describe, “the walktrap algorithm identifies multiple potential community structures based on a random series of walks (steps). Each step partitions the graph into two separate communities, merging communities in which the distance between the two communities is small enough” (6). The walktrap approach is ideal for large, directed networks such as Abraxas. The modularity score Q will be used to determine the goodness of fit of the community structure produced by the walktrap community detection algorithm. A community is typically construed as a contingent of nodes in a network that are densely connected to one another than they are to other nodes in the network. “Modularity is a chance-corrected statistic ranging from -0.5 to 1. It is defined as the fraction of ties that fall within the given groups minus the expected such fraction if ties were distributed at random” (Blondel, Guillaume, Lambiotte, and Lefebvre, 2008, 43).

According to Duxbury and Haynie (2018), modularity is calculated as:

$$Q = \sum (e_{bd} - a_b^2)$$

“where e is the fraction of ties connecting community b and community d , and a is the fraction of ties connected to community b ” (930). The higher the modularity score the more segmented a network is. Moreover, values greater than 0.3 indicate a significant community structure.

Variables and Model Estimations

To answer the third research question, three regression models were designed. In all models, the same explanatory and control variables were used with one exception. In the model which evaluated cumulative revenue generated, cumulative purchase price was not included as an explanatory variable as it was also the dependent variable.

To measure vendor trustworthiness, three proxy variables were created: success, popularity, and affluence. As trust is manifested in a variety of ways, each of these dependent variables reflects a key element of trust. Success is operationalized as the total number of transactions completed by a vendor (i.e. the number of sales made). The number of sales a

vendor makes reflects the consistent quality of their service. As an ongoing pact between the trustor and the trustee, trust is created and maintained through consistent professionalism on the part of both parties (Gambetta, 2003; Przepiorka, Norbutas, and Corten, 2017). As such, the more sales a vendor makes (with new and returning buyers), the more it is assumed that a vendor is trusted by buyers who have made an initial purchase and may return for subsequent purchases. Popularity is operationalized as the total number of unique buyers a vendor has done business with. The size of a vendor's clientele list is indicative of a more broad-based form of trust. Affluence is operationalized as the total profit a vendor has made throughout their tenure on Abraxas. This is calculated by adding the purchase price (measured in USD) of each transaction a vendor has successfully completed. Trust, in this case, is established through financial gain where reputable vendors stand to profit from the confidence buyers have in their services. Together, these dependent variables offer three distinct, though interrelated, proxies for trust. Moreover, three regression models permit for a cross-comparison of each explanatory variable's ability to explain the variance in vendor trustworthiness.

14 explanatory variables were designed (see table 1). Each reflects a measurable concept discussed within the scholarly literature regarding cryptomarket vendors (Christin, 2013; Decary-Hetu, 2016; Przepiorka, Norbutas, and Corten, 2017; Norbutas, Ruitera, and Corten; 2020). These explanatory variables are broken down into six concepts: reputation, affordability, product diversity, openness, risk-taking, and accessibility. Each of these concepts, in one form or another, help to explain vendor favourability.

Pursuant to the concept of reputation, the first explanatory variable is the cumulative reputation score. Following Decary-Hetu and Quessy-Dore (2017), the cumulative reputation score is calculated by adding the ratings of all recorded transactions a vendor has completed. Affordability reflects the costliness of a vendor. As with sellers in licit markets, cryptomarkets vendors must price their items at a reasonable rate so as to encourage buyers to do business with them. Affordability is operationalized through two variables, cumulative purchase price and average purchase price. As an aside, the cost of the product at the point of purchase is a more accurate estimate than the price as listed by the vendor. Nevertheless, the cumulative purchase price is calculated by summing all purchase prices for every transaction a vendor completes. The average purchase price is merely the average price a vendor sells a product at the point of purchase.

Product diversity reflects the variety of unique items a vendor is able to offer their buyers. This explanatory variable will implicitly contrast the profitability of product specialization with the profitability of product diversification. Indeed, the role of specialization and diversification in explaining vendor trustworthiness is yet to be understood. The concept of product diversity is operationalized through three variables, the number of unique product listings, the number of product categories, and the number of product subcategories. Each of these variables is calculated by summing up the total number of unique items or item categories within each respective category. Reflecting the concept of information asymmetry (Akerlof, 1970), openness reflects the extent to which vendors disclose product information within a listing. To clarify, each listing contains a section where vendors can provide as much or as little information on the product being sold. Openness is operationalized through the cumulative number of words variable. This reflects the number of words provided by the vendor in the description section of the listing. The cumulative number of words was calculated by summing all words in the listing for every transaction a vendor completes.

Table 1: Descriptive statistics of variables used in analysis

Variable Name	Mean or Total	SD	Median	Range
<i>Dependent Variables</i>				
Number of Transactions	20.2	38.95	7	1-330
Number of Unique Buyers	14.64	23.24	6	1-179
Cumulative Revenue Generated	2210.10	5931.95	473.25	0.23-68812.96
<i>Reputation, Price, and Risk</i>				
Cumulative Reputation	98.76	191.46	35	0-1628
Average Purchase Price	105.33	165.72	66.98	0.23-2025.04
Cumulative Risk Score	42.9	92.41	11	1-929
<i>Items and Information</i>				
Unique Items Listings	5.49	7.42	3	1-58
Unique Item Categories	1.1	0.46	1	1-5
Unique Item Subcategories	1.12	0.38	1	1-4
Number of words in item description	2773	7468.18	592	0-73267
<i>Location Shipped From</i>				
Domestic only	1700 (31.3%)	-	-	-
Regional/Continental	893 (16.4%)	-	-	-
Worldwide	2374 (43.7%)	-	-	-
Unknown	467 (8.6%)	-	-	-

Shipping internationally is generally considered a perilous activity as it increases the risk of detection when drugs move across international borders. Indeed, Branwen (2015) found that as of May 2015, 62% of cryptomarket vendors that had been arrested were apprehended in connection to international shipments. As such, risk-taking reflects a vendor's willingness to ship overseas. Risk-taking is operationalized through a cumulative risk score. As each transaction contains the locations a vendor is willing to ship to, a risk score was allocated to each transaction. However, to reduce the number of control variables, shipping locations were first pooled and set into four dummy variables to account for the different shipping categories. Risk scores were then given to each category; Unknown or N/A = missing, Domestic Only = 1 (low risk), Continental/Regional = 2 (medium risk), and Worldwide = 3 (high risk). The cumulative risk score was calculated by summing all risk scores for every transaction a vendor completes.

The final explanatory variable, accessibility, is tied to risk-taking as it refers to the locations that a vendor is willing to ship to. The more locations a vendor is willing to ship to, the less exclusive and more accessible their services are to a larger clientele base. Unlike risk-taking, the shipped to locations variable is categorical. However, as with risk-taking, shipping locations were pooled into four dummy variables to account for the different shipping categories. These include Domestic Only, Continental/Regional, Worldwide with Exceptions, and Worldwide. Importantly, Domestic Only was set as the reference category.

Trajectory Modelling

Finally, I employ k-means longitudinal modelling to determine the developmental trajectory of vendors operating on Abraxas. Like, group-based trajectory modelling (GBTM) (Nagin and Land, 1993), k-means longitudinal examines homogenous trajectories by grouping data into distinct subgroups. A hill-climbing algorithm, k-means belongs to the

Expectation-Maximization class. As such, the algorithm assigns data points to a specific cluster at the outset then recomputes each cluster to ensure that each data point moves closer to the cluster to which it best fits (Genolini and Falissard, 2010). As such, “expectation” involves a determination of the centre of each cluster while “maximization” consists of assigning each observation to the nearest cluster. These two phases are repeated until no further changes occur in the clusters.

All trajectory models were constructed in R using the KmL package (Genolini et al., 2010). Importantly, to overcome the issue of knowing a priori the exact number of clusters (or in this case trajectories) for which to group my data, I employ the Calinski Criterion to determine the optimal number of trajectory groups for each proxy variable. According to Andresen, Curman, and Linning (2016), “the Calinski Criterion is a relative metric that compares the different group solutions” (434). Importantly, a trajectory model was designed for each of the aforementioned proxy variables for vendor trustworthiness (e.g. success (completed transactions), popularity (unique buyers), and affluence (revenue)).

Findings

Descriptive Statistics

Table 2 presents a complete array of descriptive statistics for the Abraxas marketplace. In terms of the most popular drugs, Abraxas is relatively similar to other cryptomarkets, such as Silk Road 1 (Aldridge & Décary-Héту, 2016a; Christin, 2012) or Agora (Van Buskirk et al., 2016). Of the listing categories, drug and digital goods account for 92.9% (5050) and 5.9% (321) of all products sold, respectively. However, when these categories are parsed further, we can see that cannabis (34.21%), stimulants (19.38%), ecstasy (13.8%), opioids (10.8%), and psychedelics (6.75%) account for the top five products sold. This pattern can be observed in the value of transactions with cannabis, stimulants, ecstasy, opioids, and psychedelics accounting for \$198,745.16, \$149,078.46, \$95,949.28, \$94,480.70, and \$19,952.46 of the revenue generated on Abraxas, respectively. All told, this cryptomarket generated \$594,517.50 over the period of study, making it a small profit-generator relative to Silk Road 1, Evolution, Alphabay, Hansa, or Wall Street.

Table 2: Descriptive Statistics on the Abraxas Cryptomarket

Descriptive Statistics	Mean (SD) or Total	Range
<i>Vendor Reputation</i>		
Cumulative Reputation	98.76 (191.46)	0-1628
Average Reputation	4.85 (0.54)	0-5
Cumulative Positive Reputation	97.43 (189.7)	0-1625
Cumulative Negative Reputation	1.327 (4.67)	0-59
<i>Ratings</i>		
0	1.4% (74)	-
1	0.4% (23)	-
2	0.2% (10)	-
3	0.5% (26)	-
4	1.1% (59)	-
5	96.5% (5242)	-

Listing Categories

Drugs	92.9% (5050)	-
Digital Goods	5.9% (321)	-
Services	0.4% (21)	-
Drug Paraphernalia	0.3% (17)	-
Other	0.3% (14)	-
Custom Listing	0.2% (11)	-

Listing Subcategories

Cannabis	34.21% (1859)	-
Stimulants	19.38% (1053)	-
Ecstasy	13.8% (750)	-
Opioids	10.8% (587)	-
Psychedelics	6.75% (367)	-
Benzos	3.7% (201)	-
N/A	2.72% (148)	-
Prescription	2.19% (119)	-
Dissociatives	1.25% (68)	-
Information	1.03% (56)	-
E-Books	0.98% (53)	-
Erotica	0.9% (49)	-
Fraud	0.59% (32)	-
Steroids	0.35% (19)	-
RCs	0.22% (12)	-
Data	0.2% (11)	-
Drugs (Cyber)	0.17% (9)	-
Hacking	0.15% (8)	-
Money	0.11% (6)	-
Weapons	0.11% (6)	-
Electronics	0.09% (5)	-
IDs and Passports	0.07% (4)	-
Other	0.06% (3)	-
Software	0.06% (3)	-
Miscellaneous	0.04% (2)	-
Security	0.04% (2)	-
Drugs Paraphernalia	0.02% (1)	-
Services	0.02% (1)	-

Purchase Price (in USD)

All Purchases	109.41 (173.51)	0.23-2800.03
<\$1	2.2% (121)	-
\$1-\$4.99	3.3% (178)	-
\$5-\$9.99	3.1% (168)	-
\$10-\$19.99	8.7% (472)	-
\$20-\$49.99	24.7% (1344)	-
\$50-\$99.99	28.2% (1532)	-
\$100-\$199.99	16.3% (884)	-
\$200-\$499.99	10.8% (589)	-
\$500-\$999.99	1.9% (201)	-
>\$1000	0.8% (44)	-

Locations Shipped From

Australia	8.74% (475)	-
Belgium	0.83% (45)	-
Belize	0.02% (1)	-
Bulgaria	0.64% (35)	-
Canada	0.61% (33)	-
China	0.02% (1)	-
Colombia	0.02% (1)	-
Czech Republic	0.09% (5)	-
Denmark	0.81% (44)	-
Europe/EU	7.19% (391)	-
France	0.74% (40)	-
Germany	25.10% (1364)	-
Hungary	0.06% (3)	-
India	0.18% (10)	-
Italy	0.99% (54)	-
Mexico	0.02% (1)	-
Netherlands	9.22% (501)	-
Norway	0.29% (16)	-
Poland	0.11% (6)	-
South Africa	0.2% (11)	-
Spain	2.37% (129)	-
Switzerland	0.39% (21)	-
UK	13.78% (749)	-
United States	19.34% (1051)	-
Unknown or N/A	8.23% (447)	-

Locations Shipped To

Australia	8.19% (445)	-
Europe	15.73% (855)	-
Europe and US	0.07% (4)	-
Europe except Italy	0.18% (10)	-
Europe except UK	0.48% (26)	-
Germany	1.23% (67)	-
Switzerland	0.13% (7)	-
UK	4.42% (240)	-
United States	17.32% (941)	-
US and Canada	0.04% (2)	-
Worldwide	36.53% (1985)	-
Worldwide with exceptions	7.16% (389)	-
Unknown or N/A	8.60% (463)	-

As it relates to pricing, 28.2%, 24.7%, and 16.3% of products sold for prices within the ranges of \$50-\$99.99, \$20-\$49.99, and \$100-\$199.99, respectively. This suggests that Abraxas buyers did not typically spend an exorbitant amount on products. Rather, the bulk of items that were purchased were moderately priced. Nevertheless, there were 44 purchases that exceeded \$1000. Pursuant to the earlier pattern, these were purchases of cannabis (18), opioids (11), ecstasy (8), and stimulants (7). As it relates to transaction ratings, the average rating was 4.85, with 96.5% of transactions being rated as a 5. While this could mean that the

vast majority of buyers are highly satisfied with the services rendered by vendors, it may also mean that the Abraxas rating system is subject to the Pollyanna principle where there exists a positivity bias. In terms of locations shipped from, Germany, the United States, the UK, the Netherlands, and Australia are the top five shipping nations, accounting for 25.1%, 19.34%, 13.78%, 9.22%, and 8.74% of nations shipped from, respectively. Furthermore, the world, the U.S., and Europe accounted for 36.52%, 17.32%, and 17.73% of locations shipped to, respectively. Importantly, this demonstrates vendors' willingness to ship indiscriminately to all locations.

The Network Structure of Abraxas

The Abraxas transactional network is comprised of 2794 unique actors spread across 5434 transactions, with 269 unique vendors and 2525 unique buyers. There are, moreover, 3935 unique dyadic pairings. Furthermore, there are no isolates within the network as each buyer was connected to a vendor. Importantly, it was not possible to identify which buyers simultaneously operated as vendors as the unique URL tags for buyers could not be matched to unique vendor IDs. For this reason, it was not possible to calculate reciprocity or transitivity measures.

Figure 1: Abraxas Transactional Network

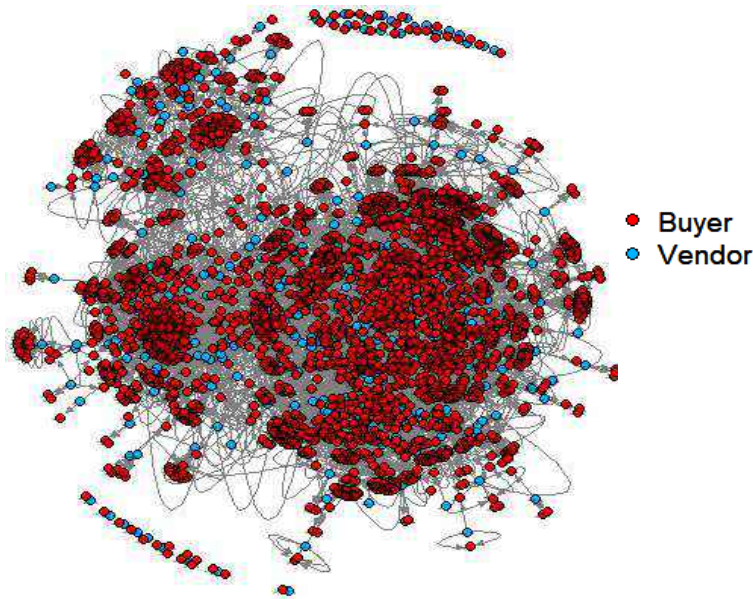


Table 3: Network characteristics

Network Characteristics	Mean (SD) or Total	Range
Unique Actors/Nodes	2794	-
Unique Vendors	269	-
Unique Buyers	2525	-
Isolates	0	-
Total Unique Edges	3935	-
Density	0.0007	-
Indegree	2.15 (2.2)	1-34
Outdegree	20.2 (39)	1-330

Indegree Centralization	0.01	-
Outdegree Centralization	0.12	-
Eccentricity (All)	11.23 (1.9)	1-16
Eccentricity (Vendors)	10.32 (3.38)	1-15
Eccentricity (Buyers)	11.33 (1.64)	1-16

The Abraxas transactional network is diffuse with a network density of 0.0007. As such, only 0.07% of all possible transactions occurred. Comparatively, Duxbury and Haynie’s (2017) cryptomarket transactional network had a density of 0.002. Furthermore, the full network consists of 29 components, with one component containing 97.6% (2726) of all nodes within the network (see table 4). The remaining connected components consisted of 19 dyads, 7 triads, and single assortments of components of various sizes. These results suggest that buyers tend to purchase from a small number of vendors over time, which leads to the formation of a large group of sparsely connected users with very few isolated buyer-seller cliques. To this extent, nodes within the Abraxas transactional network, based on the eccentricity measurement, have a maximum distance of 11.23 from one another, on average. Comparable mean values can also be observed for vendors (10.32) and buyers (11.33).

Table 4: Distribution of Network Components

Component Size	Frequency	Percentage	Node Total	Percentage
2	19	66%	38	1.4%
3	7	24%	21	0.8%
4	1	3%	4	0.1%
5	1	3%	5	0.2%
1000+	1	3%	2726	97.6%
Total	29	100%	2794	100%

Given the low network density of Abraxas, buyers did not engage with multiple vendors, doing business only with one or two with whom they trusted or were comfortable with. Indeed, 34.1% (860) of buyers purchased from two vendors exclusively while 67.5% (1702) of buyers purchased from one vendor exclusively (see table 5). As we can see, buyers prefer to do business with a small contingent of vendors as opposed to a variety. This particular preference leads to a market imbalance where a small number of vendors accounted for the majority of transactions. This can also be gleaned from the distribution of out and in-degree centrality where buyers did business with 2.15 vendors, on average, while vendors had 20.2 buyers, on average (see table 6). These findings reflect those made by Duxbury and Haynie (2017) and Norbutas (2018).

		Transactions Per Buyer												
Unique Vendors Purchased From		1	2	3	4	5	6	7	8	9	10-14	15-19	20+	Total
	1	1350	249	59	18	15	3	1	3	0	3	1	0	1702
	2	0	313	107	45	15	11	7	2	3	5	0	0	508
	3	0	0	79	50	17	11	5	4	2	3	0	0	171
	4	0	0	0	36	21	7	11	0	4	3	0	0	82
	5	0	0	0	0	9	5	5	3	3	5	0	0	30
	6	0	0	0	0	0	3	7	4	3	0	1	3	21
	7	0	0	0	0	0	0	0	1	0	0	1	1	3
	8	0	0	0	0	0	0	0	0	1	0	1	0	2
	9	0	0	0	0	0	0	0	0	0	3	0	0	3
	10	0	0	0	0	0	0	0	0	0	1	0	1	2
	11+	0	0	0	0	0	0	0	0	0	0	0	1	1
	Total	1350	562	245	149	77	40	36	17	16	23	4	6	2525

Table 5: Frequency of Unique Vendors Purchased from by Number of Transactions

However, a clearer rendering of the distribution of the in and out-degree centrality can be observed in table 6. Indeed, a majority of buyers (53.47%) transacted with only one vendor. A such, while transactions on Abraxas are not necessarily between a single buyer and vendor (i.e. there are 19 dyads), buyers generally prefer to do business with one vendor. Moreover, 22.6% of buyers transacted with two vendors. This selectivity is, understandably, not present among vendors where 84.4% have more than one buyer. Indeed, vendors will do business with a variety of buyers.

Table 6: Distribution of In and Out Degree

Degree Centrality	Outdegree Total (Vendor)	Indegree Total (Buyer)
1	42 (15.6%)	1350 (53.47%)
2	30 (11.2%)	562 (22.26%)
3	21 (7.8%)	245 (9.7%)
4	19 (7.1%)	149 (5.9%)
5	8 (3%)	77 (3.05%)
6	11 (4.1%)	40 (1.58%)
7	7 (2.6%)	36 (1.43%)
8	10 (3.7%)	17 (0.67%)
9	7 (2.6%)	16 (0.63%)
10-14	27 (10%)	23 (0.91%)
15-19	18 (6.7%)	4 (0.16%)
20-29	15 (5.6%)	5 (0.2%)
30-49	25 (9.3%)	1 (0.04%)

50-99	22 (8.2%)	-
100+	7 (2.6%)	-
Total	269 (100%)	2525 (100%)

Abraxas possesses an out-degree centralization of 0.12. Again, this is indicative of the fact that the majority of buyers typically did business with only a small number of very influential vendors. Nevertheless, some buyers were more enthusiastic in their purchasing habits than others. Whereas the average buyer made purchases from just two vendors, the most enthusiastic buyers have made purchases from over 30 vendors (range: 1-34). As it pertains to the indegree centralization of Abraxas (0.001), most buyers did not purchase very often. It is difficult to determine why this is as a buyer might have a myriad of reasons for their particular purchasing pattern. It is possible that these buyers migrated to another cryptomarket or stopped operating on the dark web altogether given the risks associated with doing so.

Interestingly, while a small number of vendors accounted for the majority of sales, the vendors outside of this power few had a difficult time earning a living on Abraxas. This can perhaps be attributed to the role of trust and reputation whereby the vendors with the best reputations continued to make sales, further increasing the barriers to entry for new vendors. To this extent, the average cumulative reputation score of a vendor is 98.76 with a standard deviation of 191.46. Moreover, these scores ranged from 0 to 1628. This is telling as the most reputable vendors attracted the most buyers, relying on their history of reputable service as a major selling point. This can be gleaned from the community analysis below.

Community Detection Analysis

Community detection analysis reveals key characteristics that provide insight into the underlying structure of the Abraxas transactional network. Abraxas possessed a total of 158 unique communities which were formed around the most popular vendors. Moreover, the community detection analysis returned a modularity score of 0.72, a relatively high Q value. This indicates that this network was heavily segmented in many communities. The largest community possessed 390 members, whereas the smallest 111 communities had fewer than 10 (see table 7). To this extent, 35 and 20 communities were dyads and triads, respectively. Indeed, the leading 20 communities accounted for 63% (1763) of all actors and 71.9% (3909) of all transactions. Moreover, the average community had 1.7 vendors and 15.98 buyers. In other words, each vendor and their respective buyers constituted an individual community.

Figure 2: Abraxas Transactional Network by Community




Table 7: Community Network Characteristics

Network Characteristics	Mean (SD)	Range
Community Size	17.7 (44.7)	2-390
Community Density	0.26 (0.19)	0.01-1
Edges	26.96 (85.81)	1-810
Within community Transactions	34.39 (103.03)	1-921
Average Cumulative Vendor Reputation	66.09 (87.97)	1-550
Avg. Outdeg (Vendor)	10.33 (12.72)	1-85
Avg. Indeg (Buyer)	1.29 (0.31)	1-2.17
Numbers of Vendors	1.7 (2.87)	1-29
Number of Buyers	15.98 (42.03)	1-373

As it relates to the composition of the communities, those with the most members had the highest average vendor reputation scores (see table 8). These communities also possessed the largest number of vendors. These communities are responsible for the lion's share of transactions made on Abraxas as a large number of buyers gravitated to a small number of trusted vendors. However, this is likely a function of the size of these communities as larger communities are comprised of more active members. In this regard, Abraxas can be characterized as a set of transactional islands that are based around several highly popular vendors who attract a large contingent of buyers. On average, the vendor-to-buyer ratio in these communities is 1:19, ranging from 1:6.5 to 1:57. Indeed, three communities are wholly dominated by a single vendor. Not surprisingly, as the size of the community increases, the network density of the community increases.

Table 8: Community Network Measures (Top 20 based on community size)

Community Size	Community Density	Edges	Within community Transactions	Cumulative Reputation (M)	Vendors	Buyers
390	0.01	810	921	266.06	17	373
337	0.01	574	748	126.69	29	308
139	0.02	331	373	153.58	12	127
129	0.01	202	247	135.78	9	120
96	0.02	151	210	166.33	6	90
91	0.02	149	176	109.5	8	83
82	0.02	117	196	294.67	3	79
58	0.03	97	105	510	1	57
53	0.03	71	111	550	1	52
52	0.02	66	89	109.75	4	48
52	0.02	65	99	246	2	50
44	0.04	85	97	121.25	4	40
38	0.04	55	71	106.67	3	35
38	0.06	80	95	237	2	36
38	0.03	45	55	251	1	37
32	0.04	36	52	82	3	29
32	0.05	53	62	102.67	3	29
32	0.04	40	64	156.5	2	30
30	0.05	40	58	72.25	4	26
30	0.05	41	74	119	3	27

Furthermore, these communities appear to be country and product-specific (see table 9). Indeed, communities, on average, had 96.7% of the items traded shipped from a single country. Moreover, these items belonged to the same item category with an average rate of 97.6%. As such, these transactional communities within Abraxas are locational and restricted to an item type. For example, a community may trade predominately in drug paraphernalia which ships exclusively from Canada. This suggests that trust, as it manifests on Abraxas, may be tied not only to a vendor’s reputation but to the country they ship from and the product(s) they sell. This reflects vendor preference. This runs counter to narratives (Barratt and Aldridge, 2016) which suggest that cryptomarkets function as a globalized transactional network. As Norbutas (2018) indicated, the structure of Abraxas’ transactional network is highly localized. However, these findings document this trend in greater detail.

Table 9: Communities by Item Categories and Country Shipped From (Top 20 based on community size)

Community size	Custom Listing	Digital Goods	Drug Paraphernalia	Drugs	Other	Services	Shipping Country 1	Shipping Country 2	Shipping Country 3	Shipping Country 4	Shipping Country 5	Shipping Country 6	Shipping Country 7
390	0%	0%	0%	100%	0%	0%	93.16%	2.71%	1.95%	1.74%	0.33%	0.11%	-
337	0%	85.45%	0%	14.55%	0%	0%	36.10%	28.74%	6.42%	5.88%	5.35%	4.95%	3.07%
139	0%	0%	0%	100%	0%	0%	92.76%	6.97%	0.27%	-	-	-	-
129	0.40%	0.27%	0%	99.33%	0%	0%	96.36%	3.24%	0.40%	-	-	-	-
96	0%	0%	0%	100%	0%	0%	68.10%	23.33%	5.71%	1.90%	0.95%	-	-
91	1.72%	0%	0%	98.28%	0%	0%	99.43%	0.57%	-	-	-	-	-
82	0%	0%	0%	100%	0%	0%	100%	-	-	-	-	-	-
58	0%	0%	0%	100%	0%	0%	100%	-	-	-	-	-	-
53	0.80%	6.97%	0.27%	90.08%	0%	1.88%	100%	-	-	-	-	-	-
52	0.11%	1.95%	0%	97.94%	0%	0%	96.63%	3.37%	-	-	-	-	-
52	0%	0%	0%	95.77%	0%	4.23%	83.84%	13.13%	1.01%	1.01%	1.01%	-	-
44	1.02%	0%	0%	98.98%	0%	0%	100%	-	-	-	-	-	-
38	0%	0%	0%	100%	0%	0%	92.96%	4.23%	2.82%	-	-	-	-
38	0%	0%	0%	100%	0%	0%	100%	-	-	-	-	-	-
38	0%	0%	0%	100%	0%	0%	100%	-	-	-	-	-	-
32	0%	0%	0%	100%	0%	0%	100%	-	-	-	-	-	-
32	0%	0%	0%	100%	0%	0%	100%	-	-	-	-	-	-
32	0%	0%	0%	96.91%	3.09%	0%	79.69%	18.75%	1.56%	-	-	-	-
30	0%	0%	0%	100%	0%	0%	65.52%	22.41%	12.07%	-	-	-	-
30	0%	0%	0%	100%	0%	0%	97.30%	2.70%	-	-	-	-	-

Regression Results and Power Few Distributions

Table 10 presents the results of the multiple linear regression models for vendor success, popularity, and affluence. In all three models, the cumulative reputation score was positive and highly statistically significant. Indeed, it appears that vendor reputation is a key predictor for trust across all three proxy variables. This finding aligns with those found in other studies (Decary-Hetu, 2016; Decary-Hetu and Quessy-Dore, 2017; Duxbury and Haynie’s, 2017), albeit on a larger scale. Additionally, it appears that cumulative risk is also a statistically significant predictor across all three models. However, the coefficient estimate, while positive for both the number of transactions and number of unique buyers, was negative for cumulative revenue generated. This is a rather curious development which cannot be easily explained without dense qualitative data. As it pertains to the number of transactions and the number of unique buyers, however, this result makes intuitive sense. Indeed, the “no risk, no reward” adage holds true on Abraxas. A willingness to incur the risks that comes with shipping overseas, and in particular worldwide, increases the number of transactions a vendor can complete and the size of their clientele base. Thus, a vendor’s success and popularity are amplified if he or she is willing and able to tap into a larger market. One might assume a logical carryover to revenue generated, but the model indicates otherwise.

Table 10: Results of Regression Models


Variable name	Number of Transactions (Success)		Number of Unique Buyers (Popularity)		Cumulative Revenue Generated (Affluence)	
	Coefficient	SE	Coefficient	SE	Coefficient	SE
Intercept	-0.79**	0.27	-0.33	1.03	2389.86***	657.76
Cumulative Reputation	0.1949***	0.0016	0.077***	0.006	37.86***	3.04
Average Purchase Price	-0.0003	0.0005	-0.0001	0.001	5.58***	1.17
Cumulative Purchase Price	0.0001**	0.00002	-0.0001	0.0001	-	-
Cumulative Risk Score	0.02***	0.003	0.059***	0.011	-35.52***	7.099
<i>Items and Information</i>						
Unique Items Listings	-0.079**	0.026	0.33***	0.098	-41.97	64.01
Item Categories	0.67*	0.29	1.298	1.098	-3777.36***	675.32
Item Subcategories	0.38***	0.11	0.831*	0.404	314.88	263.70
Number of words in item description	0.00004*	0.00002	0.0001	0.0001	0.18***	0.044
<i>Shipped to locations</i>						
Continent/Region	0.118	0.2625	0.79	0.991	539.78	646.98
Worldwide	-0.228	0.1986	0.0022	0.75	592.76	488.63
AIC	832.8	-	1496.88	-	4737.46	-
BIC	878.5	-	1542.66	-	4779.71	-

AIC = Akaike Information Criteria; BIC = Bayesian Information Criteria

$p < 0.05$ * $p < 0.01$ ** $p < 0.001$ ***

Nevertheless, each model differs in what particular estimates explain the variance in vendor success, popularity, and affluence. As it pertains to the success of a vendor, the cumulative purchase price, item categories, and item subcategories are also positive predictors. While the effect of cumulative purchase price on a vendor's success is negligible, the more categorical and sub-categorical items they can offer customers increases their likelihood of success. As it pertains to vendor popularity, unique item listings and subcategories are also positive predictors. This makes intuitive sense as the more diverse a vendor's product portfolio is, the more likely he or she is to attract a larger cohort of buyers with differing purchase interests. Finally, the average purchase price and number of words in the item description are the only positive predictors of vendor affluence. This makes sense on some level as the higher the average price of a product, the more revenue a vendor stands to generate. Moreover, given this degree of exorbitant pricing, a vendor must assure the buyer that the product they are purchasing is of the highest quality. Hence, product descriptions will contain more words, reducing information asymmetry (Akerlof, 1970).

Figure 3 presents the power law distributions of vendor success, popularity and affluence. Abraxas is subject to a power law where a small number of vendors account for the majority of transactions, unique buyers, and revenue generated. To this extent, 9.3% of vendors accounted for 50% of completed transactions, 10% of vendors accounted for 47% of unique buyers, and 5.2% of vendors accounted for 50.1% of revenue generated. Indeed, Abraxas, much a like many natural (Zipf, 1949; Simon, 1955; Eck et al., 2007) and criminological phenomenon (Sherman, 2007), is subject to the whims of a power few. This high level of preferential attachment points to the importance of trust in the structure of Abraxas' transactional network.



Trajectory Results

Table 11 presents the results of the k-means trajectory models. This table shows the three proxy variables, the number of trajectories in each model, the level of each trajectory relative to the specific variable, its base crime count in January (the month of the first transaction), the trend, and the percentage of vendors within each trajectory group. These trends are defined by regression analyses of the vendors over time within each trajectory group. A trajectory is stable if the slope parameter is between -0.2 and 0.2, decreasing if below -0.2, and increasing if above 0.2 (Curman et al., 2017).

Based on the Calinski Criterion score, I identified an optimal k-means partition of four groups for models measuring success, popularity, and affluence. Importantly, the first trajectory in each model was comprised of more than 80% of all vendors on Abraxas. This indicates that the overwhelming majority of vendors failed to conduct many transactions, engage with many vendors, or generate substantial revenue during their tenure on the market. In other words, most vendors were inconsequential in driving market activity on Abraxas, failing to generate growth. Similarly, the second trajectories in each model revealed that moderately successful, popular, and affluent vendors, based on the trend of the trajectories, grew stably across each these categories but did not ultimately become highly successful, popular, and/or affluent vendors. These vendors did not break into the upper echelon of high-performing vendors on the market. Finally, based on the third and fourth trajectories in each model, the most successful, popular, and affluent vendors continued to trend in this direction until the closure of Abraxas. These vendors grew to extreme prominence within the market and remained dominant throughout their tenure on the market.

Table 11: Summary of k-means trajectories

Variable	Trajectory	Level	Base, January	Trend	% of Vendors
Number of Transactions (Success)	1	Low	0	Increasing	83.3%
	2	Moderate	0.07	Increasing	15.6%
	3	High	0	Increasing	0.7%
	4	High	0	Increasing	0.4%
Number of Unique Buyers (Popularity)	1	Low	0	Increasing	82.2%
	2	Moderate	0.07	Increasing	16%
	3	High	0	Increasing	1.1%
	4	High	0	Increasing	0.7%
Cumulative Revenue Generated (Affluence)	1	Low	0.3	Increasing	90.3%
	2	Moderate	1.4	Increasing	8.6%
	3	High	0	Increasing	0.7%
	4	High	0	Increasing	0.4%

Figure 4 plots the trajectories of each model over Abraxas' operational timeline. Each line represents the result of the regression, showing average values. In each model, trajectories three and four exhibit large increases as a small number of vendors became the most successful, popular, and affluent in a relatively short period of time. Curiously, these vendors were relatively inactive in the first two months, springing to prominence in April and

growing exponentially in the months following. A similar pattern can be gleaned from both the revenue and affluence models. To elaborate, trajectory 4 in the success model had an average of 0 transactions in the months of January and February and 3 in March, but rose to 41, 108, and 129 in April, May, and June, respectively. Similarly, trajectory 4 in the revenue model had an average cumulative revenue of \$0 USD in January and February, but rocketed to \$17,865.2, \$30,276.7, and \$18024.6 in April, May, and June, respectively. Finally, trajectory 4 of the popularity model had, on average, 0 unique buyers in January and February, but increased to 68, 80.5 and 60.5 in April, May, and June, respectively. Curiously, nearly all trajectories in each model begin to decline following May. While market competition and dark web volatility are likely explanations, it is unclear why this is the case.

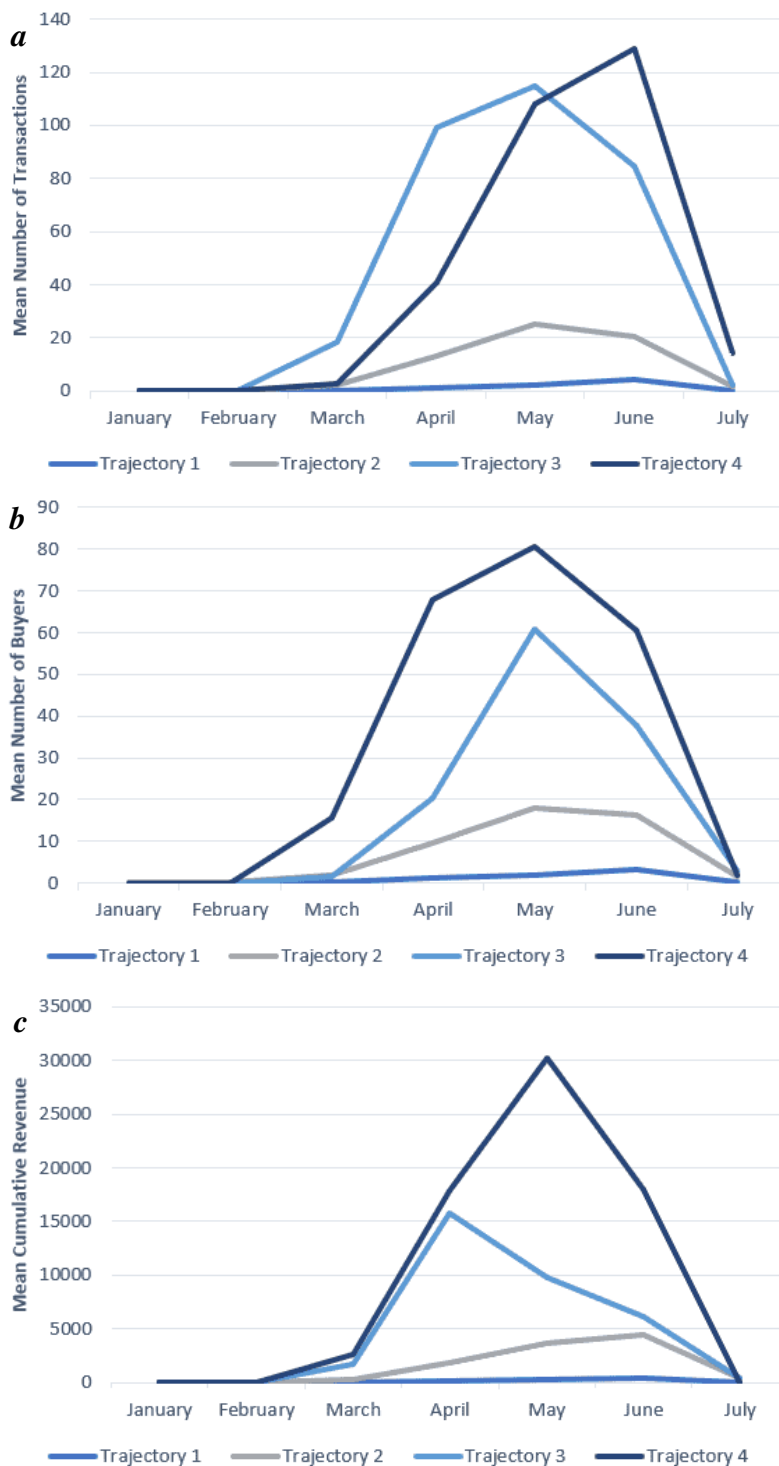


Figure 4: K-means trajectories, **a** number of transactions (success), **b** number of unique buyers (popularity), **c** cumulative revenue (affluence)

Discussion

The analyses of the Abraxas cryptomarket reveals a large and diffuse network where the majority of buyers purchase from a small cohort of vendors. This can be gleaned from the distribution of out and in-degree centrality where buyers did business with 2.15 vendors, on average, and vendors had 20.2 buyers, on average. To this extent, Abraxas is dominated by a power few of vendors that account for the majority of completed transactions, unique vendors, and revenue generated. This has important implications with regard to the development of trust in a cryptomarket as vendors who are able to create a reputation for trustworthy behaviour are most likely to succeed. Furthermore, this reputation carries over into future transactions where new buyers will do business with most trusted vendors. This can be gleaned from the results of the trajectory models where vendors with low and moderate levels of success, popularity, and affluence fail to move up while vendors high in each of these categories grow exponentially in a short period of time. Much of these findings cohere with those made by Duxbury and Haynie (2017) and Norbutas (2018). As such, these studies together shed light on the structure and trust dynamics which undergird the transactional network of cryptomarkets.

Trust, Reputation, and Network Structure on Abraxas

The network structure of an illicit market is often indicative of underlying trust dynamics (Morselli et al, 2007; Wood, 2017). Moreover, the allocation of trust within an illicit market is a paradoxical development given the high level of uncertainty therein (Kollock 1999; Yamagishi and Matsuda 2003). Relatively unexamined in the cryptomarket literature, trust dynamics are a pivotal feature which undergird market dynamics and structure. Though several studies (Duxbury and Haynie, 2017; Lacson and Jones, 2016; Janetos and Tilly, 2017) have examined and speculated on how trust is allocated in cryptomarkets, this dynamic might be better understood from the vantage point of a transactional network where vendor-buyer relations can be effectively quantified over an extended period of time. Moreover, this can be buttressed by statistical and trajectory models.

Based my findings, reputation, and to an extent, risk taking, are constituent factors which determines the network structure of Abraxas. Indeed, the power few analysis and in-degree centrality distribution suggests that a small number of vendors generate much of the market activity while buyers are inclined to do business with only these vendors. As such, the global network structure is a by-product of initial and repeated transactions between buyers and vendors. This is, moreover, meted out in the local network structure of this cryptomarket. Indeed, each vendor and their respective buyers constitute an individual community within Abraxas. These communities were also locational and product-specific, suggesting the importance of geographic distance and niche markets in moulding the network structure. To this extent, on average, 96.7% of the items traded within a community were shipped from a single country. Furthermore, these items belonged to the same item category with an average rate of 97.6%.

As such, the transactional communities within Abraxas are locational and restricted to a product type. This is contrary to narratives (Barratt and Aldridge, 2016) which suggest that cryptomarkets function as a globalized trade network where buyers and vendors from different countries engage over a variety of goods and services. Fundamentally, trust, as it

manifests on Abraxas, may be tied not only to a vendor's reputation but to the country they ship from and the product(s) they sell. However, this may reflect buyer preferences than it does vendor trustworthiness. Indeed, buyers may prefer transacting with vendors who sell a specific product and ship from a specific country due to personal preference or convenience. This is particularly important as it is often forgotten that illicit transactions are also premised on the specific desires of buyers rather than merely the trust they put in vendors. While this chapter has focused primarily on trust dynamics, this cannot be removed from buyer preferences.

How is Trust Allocated on Abraxas?

Importantly, there is a suffusion or concentration of trust, on the part of buyers, in a small number of vendors. As such, while the data cannot tell us whether or not trust is a finite commodity within Abraxas, there is evidence to suggest that it is subject to a pareto distribution. Nevertheless, buyers rely on information about sellers' past behaviour when choosing a seller. This information is based on feedback voluntarily provided by previous customers. Sellers who enter the market and have not yet established a record of good conduct can build their reputation by giving price discounts to buyers. With an increasing number of positive ratings, they can thus compensate for their initial investment by demanding a premium for their good reputation.

To this extent, trust, an artefact of the established feedback and reputation system, functions as a coordination tool on Abraxas. A buyer's feedback signals their trust, or lack thereof, in a vendor. This information can then be viewed by future buyers looking to determine the trustworthiness of that vendor. Akerlof (1970) was among the first to point out that markets run the risk of failure if buyers are unable to inspect products pre-purchase and remain uncertain as to the products' quality. Buyers having bad experiences with low-quality sellers decreases their quality expectations and thus their willingness to pay what high-quality products cost. Shapiro (1983) suggested that in order to overcome the trade-impeding information gap between buyers and sellers, high-quality sellers must invest in reputation when entering the market.

Based on the results of the regression models, vendor reputations function as a brand name, indicating to a buyer the trustworthiness and quality of a vendor. For Akerlof (1970), the deleterious effects of a lemon market can be mitigated if a buyer is able to identify the quality of merchandise. On Abraxas, reputations serve as this all-important tool for identifying the quality of merchandise and, to an extent, counteract uncertainty within a highly volatile environment. In this case, current and future buyers will then refuse to make future purchases from a low-quality vendor.

In short, reputation scores predict consumer behaviour. It is worth noting that while reputation scores reflect a seller's performance and reliability in general; it is conceivable that even high quality or reputable sellers could occasionally mislead buyers by exaggerating the quality or mislabelling a product. With that being said, a buyer might feel more anxious when considering a vendor who has a less established reputation because the vast majority of that vendor's reputation will hinge upon a small number of completed transactions. This is the opposite for vendors with many completed transactions, as we have seen. As such, it may very well be the case that there is an acceptable threshold of completed transactions for

mitigating information asymmetry and alleviating a buyer's concerns. That is, buyers may be looking for a certain amount of data on a vendor's transactional history in order to make a purchasing decision. This is particularly important as it gives us some idea about how information asymmetry is mitigated in cryptomarkets like Abraxas. Whether positive or negative, the more feedback a vendor receives the more of a known entity they become. As a result, the network structure of Abraxas may be a by-product of this dynamic.

Finally, based on the results of the trajectory models, a small number of vendors become highly successful, popular, and affluent in a relative short period of time. This is perhaps an artifact of how trust is created and distributed across a cryptomarket. As previously mentioned, while we cannot determine if trust is a finite commodity on cryptomarkets, it is unevenly distributed in a small number of vendors who reap the rewards. Moreover, this trust or lack thereof carries forward. In this case, it seems likely that trust on Abraxas is predicated on a "winner-take-all" schema where select vendors who are able to attain the trust of buyers come to dominate the market throughout its operation. Functionally, the majority of vendors who cannot establish rapport with buyers will not engage in many transactions much less generate much revenue. As such, once trust is allocated to specific vendors, it is difficult for new vendors to unseat them. In this sense, trust can be viewed as moat, functioning as a competitive advantage that separates power few vendors from the rest of those on the market.

Furthermore, what these trajectory models demonstrate is that the top vendors on Abraxas were not present or active at the inception of the market, but nevertheless came to dominate the market once they began engaging with buyers. This is perhaps reflective of a transactional cascade of sorts. Indeed, once specific vendors begin to operate on a market, completing transactions with new buyers, their activity quickly escalates, accounting for a large proportion of market activity in a fairly short period of time. It is, nevertheless, unclear whether these specific vendors were top-performing sellers on other markets that migrated to Abraxas or if they were based primarily on this market. As such, it cannot be determined if their success was organically developed on Abraxas or transferred from another market. In contrast, the vendors in the first trajectories did not see much growth across each proxy variable across time. Moreover, these vendors were present at the inception of Abraxas, making a small number of transactions in January. This indicates that a first mover principle is not present on Abraxas where early entrants to the market come to dominate market activity down the line.

Conclusion

Trust is a constitutive element of any network which trucks, barter, and trades in goods and services, regardless of their legality. In the case of cryptomarkets like Abraxas and the one examined by Duxbury and Haynie (2017), the network topology of these illicit entities are predicated on the trust buyers choose to put in the vendors they do business with. Importantly, while trust allows these transactional networks to operate in a fairly smooth manner, it may also serve to disrupt the networks' ease of operation. In this case, trust is an exploitable element within the Abraxas transactional network. Were law enforcement to design a strategy to disrupt trade on Abraxas, they would perhaps be inclined to target the vendors who were the most trustworthy in the eyes of buyers. Indeed, it is this power few

who were responsible for driving market activity on Abraxas. As a consequence, it is also likely the case that the removal of these actors would bring market activity to a halt or at the very least slow it by some degree.

In this regard, the practical implications of this study are evident. For law enforcement to effectively curtail these illicit entities, they should first begin by attempting to understand the underlying trust dynamics therein. This involves identifying the vendors that are most responsible for market activity. From here, a law enforcement organization might create a list of suitable targets for apprehension. This strategy aims to destabilize a criminal network by targeting the most trusted actors. The removal of these actors would theoretically starve a transactional network of its most pivotal economic assets, forcing buyers to switch to an unfamiliar vendor or drop out of the market altogether.

While this particular strategy makes intuitive sense based on the findings of this study, what is not well understood are the negative ramifications of targeted interventions. How would the removal of a trustworthy vendor affect the overall level of trust on the market? How might we measure this? Would buyers select another vendor on that market to do business with or would they migrate to a new market altogether? These are questions which should be examined in future studies examining the network structure of cryptomarkets. In general, cryptomarket scholars are uncertain about the impact of targeted removals. As such, scholars might run experiments on targeted interventions, testing the impact of the removal of cryptomarket vendors with real world data.

Chapter 3: Consumer Satisfaction and Information Asymmetry on a Dark Web Cryptomarket: A Text Mining Approach

For licit businesses, online reviews represent a standard metric by which consumer satisfaction and dissatisfaction are measured. Be it tourism, hospitality, or dining, numerous industries are subject to the reviews, recommendations, and opinions of their customers. This is defined as “all the informal communication directed at customers through Internet-based technology that is related to the usage or characteristics of special products and services or their providers” (Litvin et al., 2008, 201). Electronic word of mouth possesses a wider reach and facilitates faster transactions relative to other mediums where reviews might be disseminated. According to Cantalops and Salvi (2014), electronic word of mouth is more effective at driving consumer demand. Importantly, these expressions of experience introduce vital information into a market, reducing information asymmetry and allowing consumers to make more informed purchasing decisions (Akerlof, 1970). Indeed, for a market to operate with some level of cohesion, information must be readily accessible to those participating in the market.

Terrestrial criminal markets are plagued by information asymmetry. Indeed, customer satisfaction is often too difficult to measure and disseminate due to the need for secrecy and anonymity within these environments (Gambetta, 2000). However, dark web cryptomarkets represent a unique permutation which bucks this trend. Aided by anonymizing technology, cryptomarkets allow buyers to share their experience with a vendor through ratings and written reviews. This combination of quantitative and qualitative review systems is assumed to reduce information asymmetry. However, dark web scholars know very little about the intricacies of consumer satisfaction and dissatisfaction as well as the resultant reduction of information asymmetry on dark web cryptomarkets.

This study seeks to identify and compare the determinants of customer satisfaction and dissatisfaction among buyers on the Abraxas cryptomarket. It, moreover, examines the lexical predictors of vendor ratings. The overall objective of this study is to determine whether the sentiment structure of qualitative reviews differs between five-star and non-five-star ratings and how this might affect information asymmetry. This study’s secondary objective is to determine the value of “early finalization” in assisting buyers in their decision-making process. I employ a combination of text mining, sentiment analysis, and machine learning (e.g. logistic lasso regression) to identify and classify written reviews produced by buyers on Abraxas. Over the past 15 years, there has been an increase in the joint use of text mining and machine learning approaches. This combination of the aforementioned methodological techniques serves as a novel approach to a topic that has not been examined by cryptomarket researchers.

Literature Review

Customer Satisfaction and Online Reviews

The concept of consumer satisfaction and dissatisfaction has been extensively covered in research examining marketing and consumer behaviour (Chow and Zhang, 2008; Pizam and Ellis, 1999). Moreover, service quality has been identified as one major proxy for measuring consumer satisfaction (Bharwani & Jauhari, 2013; Torres & Kline, 2013). To this extent, research by Pizam and Ellis (1999), Ekinci, Dawes, and Massey (2008), and Prentice (2013) has demonstrated that tangible and intangible factors of service quality are heavily tied to how consumers rate their experience with a good or service provider. Within the context of the hotel industry, Berenzina, Bilgihan, Cobanoglu, and Okumus (2015) state, “the intangible elements are service related such as assurance, customer service and empathy whereas tangible elements are related to the physical facilities of the hotel such as appearance of hotel personnel and cleanliness of the room” (15). As such, intangible factors relate to sentimental appeals whereas tangible factors are physical characteristics. Indeed, consumer satisfaction can be achieved, maintained, and lost based on a variety of metrics relating to the performance of a business. Some measurements in combination may sour a customer’s experience while single measurements are equally likely to do so (Wilkins, Merrilees, and Herington, 2007).

The application of text mining to marketing applications is rather novel. Nevertheless, several studies have successfully employed this methodological technique to uncover hidden trends in textual feedback data. At a base level, Lee and Bradlow (2011) used text mining to examine the structure of a market based on the attributes customers had mentioned in product reviews. Ghose et al. (2011) leveraged crowdsourcing and text mining to estimate hotel demands while Archak et al. (2011) discerned patterns between the sale of electronics and the attributes listed by customers that has purchased them. Similarly, Decker and Trusov (2010) estimated consumer preferences for specific products based on the attributes they had left in their online reviews. Unsurprisingly, satisfaction with a product is correlated with its quality. However, text mining has also been applied to both measure box office performance and predict the stock performance of firms. In Eliashberg et al.’s (2007) study, the authors examined the verbiage and sentiment of movie scripts to predict their box office performance. Furthermore, Seshadi and Tellis (2012) found that “chatter” among investors, measured by the magnitude, sentiment, and rating of product reviews, determined the stock price of several companies.

Investigating the main themes motivating guests to evaluate hotels on Web 2.0, Barreda and Bilgihan (2013) determined that the cleanliness of a hotel was a primary concern of guests and often determined their level of satisfaction. Moreover, guests were more likely to produce a positively worded review for a hotel if it was located within a short distance to other venues such as shopping malls, restaurants, and airports. Finally, Pekar and Ou (2008), utilizing sentiment analysis to evaluate hotel reviews posted on “epinions.com”, found that the quality of the amenities such as food, room service, and price offered by the hotel determined the satisfaction of guests. Gan, Ferns, Yu, and Jin (2017) maintain that “star ratings in consumer-generated online reviews play an essential role in building consumer trust and are an important determinant of online business success.” Of course, the volume of

online reviews often determines the overall level of satisfaction that is expressed on the part consumers for a business.

Based on a study by Zhang et al. (2010), there is a positive association between the number of online reviews received by a restaurant and its online popularity. Thus, it may be the case that consumers tend to follow the predominant opinions of the groups or that the volume of reviews is reflective of the overall number of customers. This would indicate that a restaurant is of high quality if it is frequented by many patrons. In other words, its popularity may reflect its quality. In support of the former contention, Park, Lee, and Han (2007), demonstrate that consumers associated large volumes of reviews with favourable opinions of a product. Indeed, popularity is based on a principle of accumulated advantage where popular products become more popular. In contrast, Godes and Silva (2012) reported that negative reviews increased as the overall number of reviews increased. This makes intuitive sense as the more exposure to the public a good or service receives the more likely it is to receive negative reviews.

Cryptomarkets and Consumer Satisfaction

As with Clearnet markets, cryptomarkets employ an evaluation system where purchases are ranked with visible comments from each buyer (Resnick and Zeckhauser, 2001; van der Heide, Johnson, and Vang, 2013). That is, vendor reputations are established by consumers who are encouraged by administrators to provide publicly available feedback on their experience with a vendor. Nevertheless, while the literature on the relationship between vendor success and buyer ratings on cryptomarkets is extensive, there is distinct lack of research on qualitative customer feedback. All major studies have examined consumer satisfaction based on the reputation score of vendors, ratings left by buyers, or discussions on forums.

Hardy and Norgaard (2016) use data from Silk Road to study the relationship between reputation and prices and show that investment in reputation provides a premium to entrepreneurs. This is in line with Bhaskar et al. (2017) who demonstrate that online black markets manage to alleviate moral hazards predominantly because negative feedback led to sales reductions. In short, providing buyers with the opportunity to both air their grievances and praise vendors with whom they approve of helps the overall health of a dark market. Finally, Armona (2017) measured the impact of informal communication (through forum discussions) in anonymous marketplaces and found evidence that as the number of messages grows product demand is growing.

Janetos and Tilly (2017) show that a mature, highly rated cryptomarket vendor charges 20% higher price than a mature low-rated vendor. In general, vendors with more reviews charge a higher price than sellers with a low number of reviews regardless of rating. As such, it is speculated that vendors with a longer and more successful transactional history are more likely to cash in on this history. In other words, reputable vendors are able to exercise their brand to make a larger profit on future transactions relative to vendors without a history of success exchanges. However, bad (i.e. low-ranked) sellers prefer to exit the market than decrease their prices in response to negative feedback.

This is similar to Batikasa and Kretschmer (2018) who, studying the Agora marketplace, found that cryptomarket vendors are more likely to exit following negative

feedback. As such, it appears that receiving negative feedback early on in a vendor's tenure can reduce their chances of continued operation on a market. Negative feedback stands out more when it is not situated among positive feedback. Once a vendor is marked early as untrustworthy it is difficult to change this as buyers will not take the risk of doing business with a vendor without a proven track record for reputable economic transactions. Furthermore, a vendor's accumulated transaction experience on the platform negatively moderates market exit as a longer transactional history is correlated with continued market participation.

Finally, Przepiorka, Norbutas, and Corten (2017) use longitudinal data from Silk Road to determine to the extent to which buyers take into account sellers' reputations when making purchasing decisions. The authors conclude that "vendors react to changes in their reputation by adjusting the prices of their goods, with well-reputed vendors reaping market benefits by increasing prices" (Przepiorka, Norbutas, and Corten, 2017, 39). The authors also found that vendors with higher ratings were more successful in selling goods. Again, the successful cryptomarket vendors are able leverage their reputation to create more transactions at higher prices in the future.

Research Questions

This paper seeks to answer three research questions:

1. Based on written reviews, what are the determinants of consumer satisfaction and dissatisfaction among buyers on Abraxas?
2. Does the sentiment structure of positive and negative reviews differ? If so, to what extent?
3. What words best predict five and non-five ratings among buyers?

Identical to licit online markets, cryptomarkets utilize written reviews and ratings. However, cryptomarket research identifying the determinants of customer satisfaction and dissatisfaction is non-existent. This is the basis for the first research question. While previous studies (Christin, 2013; Decary-Hetu, 2016; Przepiorka, Norbutas, and Corten, 2017; Norbutas, Ruitera, and Corten, 2020) have examined the impact of dark market rating systems on the vendor success and profitability, none have examined this phenomenon using textual data. To clarify, while cryptomarket researchers know which vendors are reputable based on aggregate ratings and total transactions, they are generally uncertain as to why this is. To this extent, the literature is bereft of studies which examine the factors which make a vendor desirable from the perspective of buyers. As such, we are generally uncertain about the specific determinants of consumer satisfaction and dissatisfaction among cryptomarket buyers.

The second research question seeks to understand whether the sentiment structure of the written reviews match the associated rating produced by buyers. Presumably, the verbiage used in written reviews will differ as a reflection of the rating score given, with higher ratings reflecting more positive sentiments in written reviews and lower ratings reflecting more negative sentiments. However, this is still unknown among dark web researchers. Importantly, this question seeks to disentangle the Pollyanna effect which has been observed on cryptomarkets (Decary-Hetu, 2016) whereby the vast majority of ratings are very high. To

this extent, I seek to identify lexical dissimilarities in all rating categories based on the sentiment in written reviews. More generally, this research question seeks to shed light on the verbiage and tone of written reviews on cryptomarkets and how they compare to ratings.

A logical follow-up to the second research question, the third research question seeks to determine the lexical predictors of five and non-five-star vendor ratings. That is to say, this question seeks to determine if vendor ratings can be predicted from the words used in written reviews. This outlines the value of written reviews in providing prospective buyers with accurate information about vendors. On the backend, this particular question addresses the phenomenon of information asymmetry whereby information regarding the quality of products and reputation of vendors is not equally distributed to all participants in a market. Indeed, the value of feedback lies in its ability to accurately convey the experience of a buyer with a vendor such that future buyers are able to use it to their benefit when transacting with the same vendor. In short, the written feedback must go some way towards justifying the rating that was given for a transaction with a vendor.

Data

Here I use a dataset of transactions from the Abraxas cryptomarket (Branwen et al., 2015). These data contain various pieces of information from each transaction, including item title, item description, vendor name, shipping details, item reviews, items sold, transaction details, and ratings. Importantly, a customer only becomes visible once they have left a feedback following a purchase. Therefore, all active buyers were observed following their first purchase. Each recorded transaction is accompanied by feedback provided by a buyer. This includes item title, item description, shipping details, and, most importantly, written reviews produced by buyers. Each review is accompanied by the date on which it was made, the original price for which the item was bought, and a 0 to 5-star rating. The dataset contained 5434 illicit transactions between 269 sellers and 2794 buyers, over a period of 7 months in 2014–2015. Importantly, of the 5434 transactions, 4998 (92%) had a written review. These written reviews come in the form of English, French, and German textual data.

Listed on a vendor's webpage, written reviews on cryptomarkets serve as documentable proof that a transaction has occurred. "Customer feedback takes a variety of forms, ranging from detailed comments about shipping times, 'stealth' measures and the perceived potency of illicit drugs, to a simple 5-star rating" (Martin, 2014a, 41). Buyers on Abraxas are permitted to make edits to their feedback. Buyers typically provide initial feedback to indicate that the product has been purchased but will return to offer their full input once they have received and tested the product. "These comments indicate the identity of the product that was sold, the price of the sale, and the purchaser's evaluation score of the sale for all vendors' active listings" (Christin, 2013, 102).

Importantly, cryptomarket vendors, like those on Abraxas, cannot alter the feedback published on their page, whether positive or negative (Martin, 2014a, 42). As such, reputations cannot be artificially inflated by self-serving vendors. What's more, these reputations are presumably up-to-date as consumers often upload feedback upon receiving their requested product (Hout and Bingham, 2013). Tzanetakis, Kamphausen, Werse, and von

Laufenberg (2016) argue that customer feedback in cryptomarkets mitigates some of the risk associated with illicit drug trading.

While all cryptomarkets are feedback-based, they may differ on policies regarding the mandatory nature of buyer feedback. That is to say, some cryptomarkets require buyers to leave feedback after every transaction while others do not. Abraxas falls into the former category, with all transactions conducted over the market's operational period being documented via buyer feedback. While feedback data would ordinarily pose a problem in many network-based cryptomarket datasets due to partial or completely anonymized buyer nicknames, Abraxas contained unique buyer profile identifiers for each feedback message, which was located in the HTML code of item pages. I used these buyer identifiers to aggregate feedback messages left by each buyer account.

Methods

The Utility of Texting Mining

A relatively recent technological development, Mikroyannidis and Theodoulidis (2006) define text mining as the act of “processing a collection of documents, or corpus, in which documents are converted into structured data, such that each document is described using a set of features called concepts to provide a holistic perspective of textual and non-textual information” (45). More generally, text mining allows for the automatic analysis of large amounts of qualitative data, a previously arduous task. Given that this study analyses 4998 customer reviews, traditional qualitative research approaches such as grounded theory or content analysis were inadequate. Text mining represented the most viable methodological option as these tasks can be achieved via computation.

In general, text mining can be separated into linguistic and non-linguistic approaches. According to Taboada et al. (2011), “linguistic techniques consider the natural language characteristics of the text in documents (e.g., syntax, grammar)” (101). In contrast, Ur-Rahman and Harding (2011) define “non-linguistic techniques view documents as a series of characters, words, sentences, and paragraphs” (78). Given the descriptive focus of this study, non-linguistic text mining approaches will be employed to calculate the frequency and proximity of words. In particular, non-linguistic text mining allows for a term frequency-based matrix to represent the data while reducing key information loss (Ur-Rahman and Harding, 2011). More importantly, it will employ three analytic strategies: descriptive text analysis, sentiment analysis, and textual feature extraction. All analyses and visualizations were conducted in R.

Variable Operationalization and Data Pre-processing

Intuitively, positive written reviews reflect customer satisfaction while negative reviews reflect dissatisfaction. Based on Venkatesh and Goyal's (2010) expectation-disconfirmation model, consumer satisfaction and dissatisfaction are reflective of congruences between a customer's expectation of the product and their actual perception once the product has been received. This reflects individual cognitive processes when

evaluating a good or service. As such, consumer satisfaction, for the purpose of this study, is defined as an event in which the consumer's perception of the good or service purchased matches or has exceeded their expectation of the good or service prior to purchase. In contrast, customer dissatisfaction is an event in which the customer's perception of the good or service upon purchase falls below their expectation prior to purchase. Compared to qualitative feedback, numerical ratings are much simpler to gauge. The higher the rating the more the customer was satisfied with the transaction. This is not imminently clear with written reviews as customer satisfaction or dissatisfaction can only be gauged once the review is read in its entirety and compared against other reviews.

As Abraxas' rating system ranges from 0 and 5, this study operationalizes negative reviews as any rating below 5. As such, dissatisfied buyers are those who have rated their experience as anything below 5. Logic dictates that a positive rating is operationalized as a 5. But why select this specific dichotomy? Cryptomarket rating systems are seemingly subject to the Pollyanna principle or a positivity bias where buyers are more likely to remember positive experiences with vendors when producing their ratings than they are negative experiences (Decary-Hetu and Quessy-Dore, 2017). As such, the overwhelming majority of vendor ratings across a number of cryptomarkets are 5's (Decary-Hetu, 2016). While several studies (Christin, 2013; Norbutas, and Corten, 2017; Norbutas, Ruitera, and Corten; 2020) have reported a high level of quality among dark market vendors, it is unlikely that the vast majority of cryptomarket transactions are perfect as stipulated by a rating of 5. As 5 ratings are the rule and not the exception, any rating below a 5 is considered an anomaly. Thus, ratings below 5 are designated as negative while ratings of 5 are designated as positive.

Importantly, text mining necessitates a series of pre-processing procedures. For this study, data pre-processing consisted of tokenization, filtering, and stemming. More specifically, the textual data were cleaned by removing punctuation, special characters, digits, and uniform resource locator links. Tokenization was then conducted. Tokenization is the process reducing words into pieces of information called tokens. The objective of tokenization is the identification of meaningful keywords. Next, all stop-words were removed from the corpus. Stop-words are functional fillers which do not carry any information. According to Liua and Tan (2017), "prepositions (such as 'from', 'to', 'after', etc.), articles (such as 'a', 'an' and 'the') and pronouns (such as 'I', 'you', 'she', 'he', etc.) can be treated as stop-words" (56).

Next, word stemming is conducted. Word stemming involves breaking words down to their roots (Liau and Tan, 2017). These data were then converted to a corpus. From here, these data were converted into a structured format from which analyses can be conducted. Finally, a vectorspace model is created. This step is required for feature extraction. "Each document is represented as a vector (v) in the (t) dimensional space if we have a set of (d) documents (i.e. written reviews) and a set of (t) terms" (Elagamy, Stanier, and Sharp, 2018). The feature extraction stage produces a two-dimensional matrix (vector space). I then produced a TF/IDF (term frequency–inverse document frequency) value for each feature. The TF/IDF is a numerical statistic which reflects a word's importance to a document in a corpus.

Descriptive Text Analysis

Descriptive text analysis is a fairly standard text mining procedure. Simple term frequencies are conducted to identify the words used by Abraxas buyers to describe their experience. This is done for the entire corpus as well as for fives and non-fives. The frequency and distribution of specific terms across the corpus provides insight into the nature of cryptomarket activity. While an examination of word frequency will often reveal words that are expected, there is always the possibility of discovering usual words that offer more insight into cryptomarket transactions. I also utilize hierarchical cluster analysis to identify the optimal number of word clusters within the word cluster. In particular, I employ agglomerative clustering to fuse individual words into groups by measuring the distance between term vectors. This particular method was employed due to its simplicity and ease of use for textual data.

Agglomerative hierarchical cluster analysis takes a table of i individuals (rows) and j variables (columns) and converts it into a distance matrix. “The analysis starts with each individual in a single cluster and then combines individuals progressively into larger clusters until a final stage where all individuals are merged into a single group” (Desagulier, 2019). A tree-like dendrogram is then used to graphically represent this stepwise process. A dendrogram is a tree-like visualization based on frequency distance of words. Importantly, to reduce the sheer number of redundant terms, I set the sparsity threshold of the term document matrix to 95%. As the sparse parameter is a number between 0 and 1, setting a sparse parameter of 95% only includes words with 95% or fewer zeros (appear only once). As most corpora are likely to have 0.95 or more zeros, setting a dendrogram based on a sparsity of 95% is methodologically prudent. This removes words that are found in only 5% of the data. In addition to hierarchical cluster analysis, I employ word associations. Similar to the statistical concept of correlation, word association measures the frequency in which words co-occur (Correia, Teodora, and Lobo, 2018).

Sentiment Analysis

I conduct sentiment scoring or polarity calculations on the written reviews. Sentiment analysis is the process of determining the positive, negative, or neutral sentiment in textual data. This, moreover, comes in the form of a score. In business, companies use sentiment analysis “to develop their strategies, to understand customers’ feelings towards products or brand how people respond to their campaigns or product launches, and why consumers do not buy some products” (D’Andrea et al., 2015, 27). The “qdap” package in R provides a polarity function which is accurate and uses basic arithmetic for scoring. This dictionary ranges from -1 to 1, with -1 and 1 reflecting negative and positive sentiment, respectively. This package features an extensive sentiment library of adjectives and phrases that were hand-scored by human coders.

Sentiment analysis will be done separately for five and non-five ratings in order to determine whether or not the sentiment of supposedly positive and negative reviews align with the verbiage used. Furthermore, sentiment analysis is conducted on feedback that are

finalized early and those that are not. In short, I seek to determine if the early finalization is correlated with a positive transactional experience on the part of vendors. In summation, sentiment analysis offers some idea about the usefulness of written reviews as a means of mitigating information asymmetry (e.g. quality of the information). As sentiment scoring measures the positive or negative intent in a buyer's tone, future buyers are able to quickly discern the experience of past buyers relative to the rating they provided on a transaction. The more detail a buyer provides in their feedback the more information a prospective buyer has on the quality and trustworthiness of a vendor.

Feature Extraction: Logistic Lasso Regression

Though exploratory text analysis offers descriptive insights into buyer reviews, it does little with regard to classification and prediction. Feature extraction is based on dimensional reduction where large datasets are made into smaller, manageable dataset through which more suitable statistical techniques can be performed. In this article, I use the written review text as predictor variables to classify whether a written review will be positive or negative (i.e. satisfied or dissatisfied). Classification is a type of machine learning exercise which predicts the most probable label Y for an instance X. In this case, feature extraction is used to understand what lexical elements (e.g. words) predict for customer satisfaction and dissatisfaction. To this effect, I will be using a supervised machine learning technique: logistic lasso regression.

Determining which textual predictors are associated with an outcome is not a simple task. In linear regression, one attempts to model a dependent variable using the best straight line fit to a set of predictor variables. Moreover, when selecting the variables for a linear model, one generally looks at individual p-values. Given that many lexical features are superfluous and the outcome variable is binary (satisfied/dissatisfied or five/non-five), a linear regression cannot be used. Instead, a logistic lasso regression will be used.

According to Park and Casella (2008), "One can think of logistic regression as the equivalent of linear regression for a classification problem. It is a regression analysis where the response variable is binary, meaning that it can only assume 0 (dissatisfied) or 1 (satisfied) values. The explanatory variables can be either discrete or continuous." The Least Absolute Shrinkage and Selection Operator or LASSO is a regularization method in statistical modelling that is used when the data is noisy and outcome variable is binary. The LASSO method puts a constraint on the sum of the absolute values of the model parameters. This method applies a shrinking (regularization) process where it penalizes the coefficients of the regression variables shrinking some of them to zero. Traditionally, one might engage in feature selection by manually evaluating the p-values of coefficients and removing those that are not statistically significant. This, however, can be a laborious process if there are swaths of coefficients. LASSO automatically selects significant coefficients by shrinking unimportant predictors to zero.

The effect of the penalty term is to set certain coefficients exactly to zero. During the feature selection process, the variables that still have a non-zero coefficient after the

shrinking process are selected to be part of the model. The goal of this process is to minimize the prediction error. In order to assess how well the model can be generalized to my dataset, I utilize k-fold cross-validation (Hastie et al., 2009) where the data are portioned into four subsets of approximately equal size and one of the subsets becomes the validation set. The remaining three subsets are used as training data. Abiding by suggestions from Pereira, Basto, and da Silva (2015), this procedure is repeated 10 times, each time with a different validation set, and the optimum value of λ is estimated such that the cross-validated log-likelihood is maximized (Goeman, 2010).

Finally, the logistic lasso regression will also provide odds ratios for all predictors. An odds ratio is a measure of association between an exposure and an outcome. “When a logistic regression is calculated, the regression coefficient (β_1) is the estimated increase in the log odds of the outcome per unit increase in the value of the exposure” (Szumilas, 2010, 227). In other words, the exponential function of the regression coefficient is the odds ratio associated with a one-unit increase in the exposure. These will be used to understand the influence of various words in buyer satisfaction and dissatisfaction on Abraxas.

Regardless, there is one glaring limitation of these feature extraction processes: class imbalance in binary classification. As there are a total of 175 non-five ratings relative to 4683 five ratings, these binary classification models will be trained on a very small number of dissatisfied customer reviews. However, class imbalance is in fact a fairly common classification problem in machine learning. While there is no perfect solution to this limitation, using a penalized model like lasso logistic regression imposes additional costs on the model for making a classification mistake on the minority class during training. While a lasso logistic regression by no mean solves the problem of class imbalance, it serves to reduce its impact.

Findings

Following the exclusion of Abraxas transactions without feedback data and removal of non-English feedback, 4858 total transactions remained. Table 1 presents descriptive statistics for all product types by customer ratings. As is imminently clear, Abraxas feedback is governed by the Pollyanna principle where 96% (4683) of all purchases received a rating of five. However, this positivity bias is not limited to product type as all categories have a mean rating above four. Moreover, given relatively similar average ratings for transactions that have been finalized early (4.83) and those that have not (4.9), early finalization does not necessarily imply greater consumer satisfaction. This perhaps suggests that early finalization is premised on expedience than it is on customer satisfaction as buyers presumably trust a vendor and wish to receive their purchase sooner. Nevertheless, it would seem from the quantitative data that Abraxas buyers are overwhelmingly satisfied with their transactions.

Table 1: Descriptive statistics (ratings by product type and finalize early)

Variable Name (N)	Mean	SD	Median	Range	Rating 0	Rating 1	Rating 2	Rating 3	Rating 4	Rating 5
All (4858)	4.89	0.67	5	0-5	67	22	10	24	52	4683
Custom Listing (11)	4.09	2.02	5	0-5	2	0	0	0	0	9
Digital Goods (253)	4.91	0.58	5	0-5	2	1	3	0	3	245
Drug Paraphernalia (16)	5	0	5	5	0	0	0	0	0	16
Drugs (4548)	4.89	0.68	5	0-5	63	21	8	24	49	4383
Services (19)	5	0	5	5	0	0	0	0	0	19
Other (11)	5	0	5	5	0	0	0	0	0	11
Finalize Early	4.83	0.83	5	0-5	24	9	2	11	8	1076
Not Finalize Early	4.9	0.62	5	0-5	43	13	8	13	44	3607

Table 2 presents the word and character counts for Abraxas feedback by product type, rating, and purchase price. Indeed, it does not appear that the length and detail of consumer feedback differs significantly based on the type of product purchased and the amount the product was purchased for. However, there are noticeable derivations within these two categories. While feedback for custom listings, drugs, and services have, on average, 11.09, 9.81, and 9.05 words, respectively, feedback for digital goods, drug paraphernalia, and other products have an average 5.15, 6.88, and 6 words, respectively. Though it is difficult to know why the length of feedback for these products differ, it may be that the experiential nature of these products and services lends themselves to differing feedback content. For example, while a buyer purchasing a hacking guide describes the transaction as “nice tutorial :)”, another buyer purchasing marijuana describes the transaction as “Trusted Vendor, was a little overweight. Not AAA Weed but very good.” Indeed, the fact that drug transactions are predicated on the weight and quality of the product while digital good transactions pertain to the simple functioning of the product, the length and detail of the feedbacks will differ. Therefore, the type of product purchased seems to pre-empt the length of the feedback

Table 2: Word and character statistics

Variable Name (N)	Mean Word Count (SD)	Range of Word Count	Mean Character Count (SD)	Range of Character
<i>Product Type</i>				
All (4858)	9.55 (8.89)	1-99	48.16 (41.45)	1-412
Custom Listing (11)	11.09 (7.18)	3-28	63 (38.9)	16-148
Digital Goods (253)	5.15 (7.07)	1-48	25.75 (32.35)	1-214
Drug Paraphernalia (16)	6.88 (5.06)	1-19	37.62 (27.22)	6-103
Drugs (4548)	9.81 (8.92)	1-99	49.48 (41.57)	1-412
Services (19)	9.05 (12.05)	1-46	42.16 (51.15)	4-209
Other (11)	6 (4.84)	2-18	30.64 (21.95)	9-69
<i>Ratings/FE</i>				
0 (67)	19 (14.52)	1-50	92.24 (66.64)	4-248
1 (22)	19.59 (13.87)	1-49	91.45 (59.01)	5-214
2 (10)	19.4 (17.21)	4-51	86.7 (75.43)	19-209
3 (24)	19.46 (14.98)	1-66	97.71 (72.05)	8-329
4 (52)	18.98 (15.4)	1-69	87.17 (68.83)	7-294
5 (4683)	9.19 (8.39)	1-99	46.56 (39.29)	1-412
Finalize Early (1130)	9.1 (8.54)	1-99	45 (39.1)	1-412
No Finalize Early (3728)	19.59 (13.89)	1-49	91.45 (59.01)	5-214
<i>Purchase Price (in USD)</i>				
<\$1 (98)	6.77 (9.32)	1-48	32.12 (39.74)	2-208
\$1-\$4.99 (143)	6.49 (7.98)	1-46	32.43 (38.82)	1-233
\$5-\$9.99 (137)	7.18 (7.38)	1-50	35.66 (32)	3-200
\$10-\$19.99 (418)	10 (9.46)	1-51	49.66 (43.67)	1-254
\$20-\$49.99 (1205)	9.1 (8.34)	1-99	45.88 (38.96)	1-412
\$50-\$99.99 (1373)	9.91 (8.82)	1-68	50.23 (41.29)	2-329
\$100-\$199.99 (795)	9.66 (8.96)	1-69	49.26 (42.17)	1-294
\$200-\$499.99 (548)	10.58 (9.47)	1-68	53.16 (43.86)	2-310
\$500-\$999.99 (99)	11.83 (10.77)	1-64	59.76 (48.7)	4-268
>\$1000 (42)	10.24 (7.94)	1-31	50.14 (36.85)	3-158

As it pertains to purchase price, it seems that the lower the purchase price the shorter the feedback. Indeed, the cheapest purchase price categories, <\$1, \$1-\$4.99, and \$5-\$9.99, have less words per feedback relative to the other purchase categories which contain a similar number of words per feedback (ranging from 9.1 words to 11.83 words). This indicates that buyers who make expensive purchases are generally more expressive than buyers who makes inexpensive purchases. This is perhaps due to the quality of the product where buyers are more inclined to describe the product in detail. Moreover, this is particularly useful for prospective buyers looking to make similar expensive purchases. Indeed, the more information that is provided on expensive products the less information asymmetry buyers need to contend with when making the decision to purchase.

Importantly, there is a substantial difference in both the average word and character count between five-rated transactions and those with non-five ratings. Indeed, feedback for

five-rated transactions possess nearly 10 less words on average than non-five-rated feedbacks. What this might suggest is that buyers that have not rated their transaction as a five are more inclined to go into detail as to why this was the case. As non-five transactions are rare on Abraxas, the resultant feedback goes beyond merely praising the vendor and/or product. Finally, I observe large differences in the average word and character count between transactions that have been finalized early and those that have not. To this extent, transactions that have been finalized early receive approximately 10 and 46 less words and characters than transactions that have not been finalized early, respectively.

Descriptive Text Analysis

Based on simple frequency analyses (see Figures 1(b) and 1(c)), there are some glaring similarities and differences between the words used to describe five-rated and non-five-rated transactions. To this extent, “finalize” and “early” are the two most popular words among non-five-rated feedback and the second and third most popular words among five-rated feedback. However, it is also evident that the tone of words differs between these two groups. Indeed, it appears that many of the most frequently used words within the five-rated corpus possess a positive connotation: “good”, “great”, “fast”, “stealth”, “thanks”, “quality”, “trust”, “best”, “top”, and “nice”. In contrast, words in the non-five-rating corpus are negative or value-neutral: “update”, “scam”, “product”, “nothing”, “never”, “still”, and “waiting”. This suggests inherent lexical dissimilarities between five and non-five feedback. Moreover, it is evident that buyers that have rated their transaction as a five are more satisfied with the product and vendor.




Figure 1(a): Word frequencies top 20 (all)




Figure 1(b): Word frequencies top 20 (<5's)

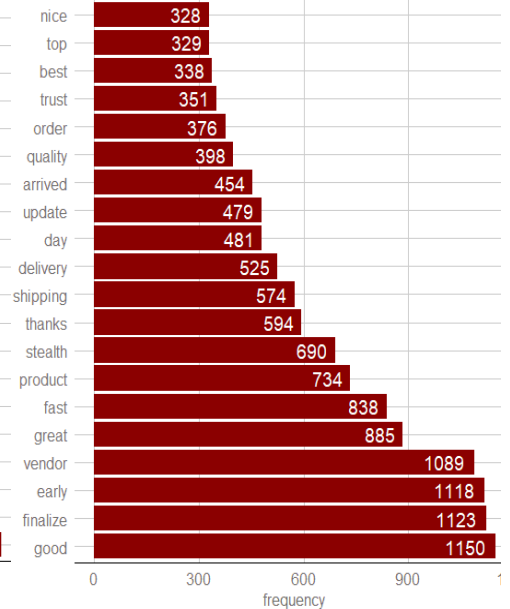


Figure 1(c): Word frequencies top 20 (5's)

Abraxas feedback is subject to Zipf's law where 1% of words occurred 47.3% of the time (see Figure 2). Moreover, 20% of all words were used in 94% of all instances. This discrete pareto distribution is not an altogether surprising as the same pattern occurs in aggregated conversations in the English language. As it relates to cryptomarket feedback, buyers will typically use the same words to describe their experience with a vendor. While this does not necessarily mean that buyers are not very expressive, it likely means that a small number of words suffice in describing a buyer's experience. This can be gleaned from table 2 where all transactions with feedback had an average of 9.55 words. This is further explored in table 3.

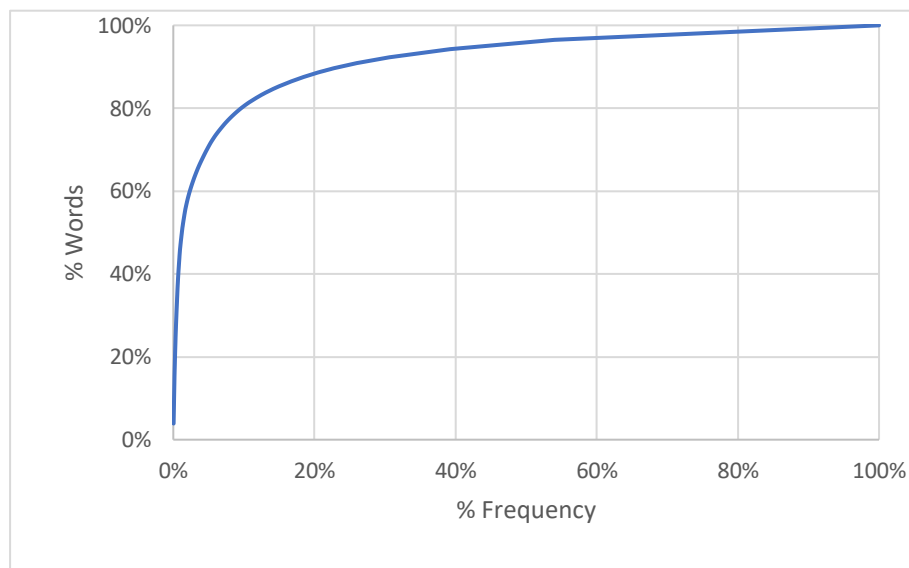


Figure 2: Power law distribution for word frequency

These observations are further present when these corpora are disaggregated to control for product type and ratings (see Table 3). With the exception of digital goods, “finalize” and “early” are among the top five most popular words within each product-based corpus. More generally, the words used in each product-based corpus are seemingly positive, suggesting that buyers were satisfied with the products purchased. While “finalize” and “early” are also present in ratings-based corpus (with the exception of 4-rated transactions), there are noticeable differences in the connotation of the words used. Indeed, higher ratings contained more positive words while lower ratings did not. In particular, “scam” was the third-most used word among transactions that were rated zero. Though used only three times, “bad” was a frequently occurring word among two-rated transactions. In contrast, “good” and “great” are frequently occurring words among transactions rated a three, four, and five.

Table 3: Word frequency top 5 by product type and rating

Variable Name	1st	2 nd	3rd	4 th	5th
<i>Product Type</i>					
Custom Listing	Early (6)	Finalize (6)	Best (3)	Day (3)	Good (3)
Digital Goods	Good (47)	Thanks (44)	Fast (43)	Vendor (41)	Great (39)
Drug Paraphernalia	Great (5)	Early (4)	Finalize (4)	Good (3)	Service (3)
Drugs	Finalize (1171)	Early (1167)	Good (1127)	Vendor (1073)	Great (843)
Services	Early (4)	Finalize (4)	Great (4)	Day (3)	Order (3)
Other	Early (3)	Finalize (3)	Day (2)	Described (2)	Easy (2)
<i>Ratings</i>					
0 (67)	Early (31)	Finalize (28)	Scam (27)	Order (20)	Update (19)
1 (22)	Order (15)	Early (10)	Finalize (10)	Vendor (10)	Day (8)
2 (10)	Bad (3)	Day (3)	Early (3)	Finalize (3)	Order (3)
3 (24)	Early (17)	Finalize (15)	Update (14)	Arrived (12)	Good (9)
4 (52)	Order (19)	Day (18)	Good (17)	Arrived (16)	Product (12)
5 (4683)	Good (1149)	Finalize (1123)	Early (1118)	Vendor (1089)	Great (885)
Finalize Early	Early (1184)	Finalize (1180)	Update (429)	Vendor (364)	Arrived (238)
No Finalize Early	Good (1033)	Great (791)	Vendor (761)	Fast (751)	Product (631)

Figures 3(a), 3(b), and 3(c) present cluster dendrograms for all corpora. The key to interpreting a dendrogram is to focus on the height at which any two objects are joined together. Moreover, the heights of joined words reflect the distance between the clusters. Among non-five-rated transactions, “finalize” and “early” are again connected. However, “nothing” and “shipping” are a predominant pairing, suggesting that dissatisfied buyers base their dissatisfaction around not receiving the promised good or service. Furthermore, “scam” is associated with this cluster, again suggesting that buyers perceive the absence of a purchased item upon delivery as a con. “Good” and “product” also form a pair suggesting that some of these buyers, despite not rating the transaction as a five, still found some utility in the transaction. As it relates to the cluster dendrogram for five-rated transactions (see Figure 3(c)), “finalize” and “early” are again a predominant pairing. Moreover, “fast” and “shipping”, “fast” and “delivery”, “fast” and “stealth” are pairings associated with “great”. This suggests that satisfied buyers are concerned with the speed and stealth of the product purchased.




Figure 3(a): Cluster dendrogram (all)




Figure 3(b): Cluster dendrogram (<5's)




Figure 3(c): Cluster dendrogram (5's)

These trends are further present when examining word associations for the five most frequently used words among five and non-five-rated transactions (see Tables 4(a) and 4(b)). In both corpora, “finalize” and “early”, while highly associated with one another, also co-occurring with “trust”, “hope”, and “confidence”. Unsurprisingly, buyers put a significant amount of trust in vendors when deciding to forgo the use of Abraxas’ escrow system and finalize their purchases early. Nevertheless, there are differences in word association in these corpora. Notably, five-rated words associations are characterized by trust, satisfaction, and praise for the vendor, product, and process. For example, “good” is associated with “stealth”, “price”, “communication”, and “product” while “great” is similarly associated with “product”, “stealth”, “communication”, “shipping”, and “vendor”. In contrast, non-five word associations are generally value-neutral but are negative when the transaction is described. As such, “order” is associated with “theft”, “risk”, “mistake”, “hostile”, and “defiant”.

Table 4(a): Word Associations (5’s)

Word Rank	Good (1182)	Finalize (1191)	Early (1188)	Vendor (1125)	Great (893)
1	Stealth (0.19)	Early (0.99)	Finalize (0.99)	Trust (0.41)	Product (0.2)
2	Price (0.15)	Update (0.47)	Update (0.47)	Best (0.16)	Stealth (0.17)
3	Communication (0.14)	Trust (0.27)	Trust (0.27)	DNM (0.12)	Communication (0.15)
4	Look (0.13)	Arrived (0.2)	Arrived (0.2)	Great (0.12)	Shipping (0.13)
5	Fast (0.12)	Upon (0.15)	Upon (0.15)	Finalize (0.11)	Service (0.12)
6	Product (0.11)	Confidence (0.13)	Confidence (0.13)	Early (0.11)	Vendor (0.12)
7	Stuff (0.11)	Vendor (0.11)	BTW (0.12)	Professional (0.1)	-
8	-	Later (0.1)	Received (0.11)	-	-
9	-	Received (0.1)	Vendor (0.11)	-	-
10	-	Hope (0.1)	Hope (0.1)	-	-

Table 4(b): Word Associations (<5’s)

Word Rank	Early (70)	Finalize (68)	Order (61)	Arrived (51)	Day (46)
1	Finalize (0.96)	Early (0.96)	Maybe (0.36)	Update (0.46)	Ago (0.41)
2	Update (0.46)	Update (0.47)	Decent (0.34)	Finalize (0.3)	Marked (0.41)
3	BTW (0.41)	Arrived (0.3)	Theft (0.31)	Upper (0.25)	Rewording (0.37)
4	Yet (0.37)	Answer (0.28)	Risk (0.31)	Quantum (0.25)	Pay (0.37)
5	Answer (0.36)	Ganja (0.28)	Mistake (0.31)	Big (0.25)	Meds (0.37)
6	Soon (0.26)	Baggy (0.28)	Spain (0.31)	Technical (0.25)	Choice (0.37)
7	Doesnt (0.26)	Trust (0.27)	Hostile (0.31)	Contact (0.25)	Accurate (0.37)
8	Ganja (0.26)	Must (0.2)	Digits (0.31)	Easy (0.25)	Weight (0.37)
9	Arrived (0.25)	Cant (0.2)	Defiant (0.31)	Early (0.25)	Later (0.34)
10	Domestic (0.2)	Anything (0.2)	Attempt (0.31)	Order (0.24)	White (0.34)

Sentiment Analysis (Ratings)

Based on the sentiment analysis, there is a clear difference between the five-rated and non-five-rated corpora (see Table 5). To this extent, the five-rated corpus has a positive average polarity of 0.6 while the non-five-rated corpus is negative with an average polarity of -0.01. Intuitively, the average polarity of the combined is 0.58, suggesting that the addition of the non-five-rated transactions slightly lowers the overall sentiment score. Importantly, the non-five-rated corpus is only slightly negative, verging on an average polarity which is neutral. This suggests that buyers, while dissatisfied, will not harshly criticize vendors or vent their frustrations when providing feedback to the rest of the market. This distribution of sentiment scores can be viewed in Figure 4(a), 4(b), and 4(c).

Table 5: Distribution of sentiment polarity (ratings)

Corpus	Total Sentences	Total Words	Average Polarity	Std Dev Polarity	Std Mean Polarity
All	4858	45103	0.58	0.56	1.03
5	4683	41849	0.6	0.56	1.08
<5	175	3254	-0.01	0.41	-0.01




Figure 4(a): Sentiment polarity (all)

Figure 4(b): Sentiment polarity (<5's)

Figure 4(c): Sentiment polarity (5's)

Figure 5 presents the most frequently occurring high-sentiment words in each corpus; five-rated transactions are in green while non-five-rated transactions are in red. This figure contrasts the words used in either category. As is evident, five-rated transactions possess a higher frequency of positive words than non-five-rated transactions. These words include “awesome”, “always”, “fantastic”, “confidence”, and “love”. To this extent, words in the non-five corpus are generally negative but verging on value-neutrality. These include “avoid”, “frustrating”, “incorrectly”, “seized”, and “caught”. These particular words indicate

that buyers were generally dissatisfied with a transaction due to it not arriving or not arriving as it was advertised. These particular buyers suggested that prospective or future buyers “avoid” doing business with this vendor. This offers some evidence as to the provision of suggestive feedback from dissatisfied buyers. In this case, these buyers are advising future buyers to do business with more trustworthy vendors.



Figure 5: Sentiment cloud; 5's and <5's

Sentiment Analysis (Finalize Early)

While the average rating did not differ between transactions that were finalized early and those that were not, the sentiment analysis reveals a clear difference between these transaction types (see table 6 and figures 6(a) and 6(b)). To this extent, the finalize early corpus had a lower average polarity (0.31) than the non-finalize early corpus (0.66). This suggests that those who finalized early were less likely to use positive verbiage in their feedback relative to those who did not finalize early. This is particularly surprising as buyers were generally less positive about transactions they had finalized early. What this might suggest is that the expectation of buyers who finalized early were not met.

This is reflective of Venkatesh and Goyal's (2010) expectation-disconfirmation model where consumer satisfaction and dissatisfaction are based on congruences between a customer's expectation of the product and their actual perception once the product has been received. To this extent, given that buyers had finalized early, it is possible that their expectations of the product were inflated relative to buyers that did not finalize early. Indeed, it is likely that these buyers bought into the hype of the product based on the vendor's reputation and the number of prior transactions they had made. As such, once the product was received their inflated expectations did not match up the actual quality of the product. This is

not observed among buyers that did not finalize early as their expectations were perhaps more in line with the actual quality of the product.

Table 6: Distribution of sentiment polarity (finalize early)

Corpus	Total Sentences	Total Words	Average Polarity	Std Dev Polarity	Std Mean Polarity
FE	1130	12183	0.31	0.41	0.75
No FE	3728	32920	0.66	0.58	1.15




Figure 6(a): Sentiment polarity (FE)

Figure 6(b): Sentiment polarity (No FE)

Feature Extraction

Figure 7 presents the error as a function of lambda of the logistic lasso regression. The plot has the mean square error on the y-axis and the natural log of λ on the x-axis. The plot shows that the log of the optimal value of lambda is 3.28, with the minimum lambda value of 0.0005 and maximum of 0.014. Importantly, the optimal lambda indicates the accuracy of the model, with a higher value equating to a more accurate the model. In this case, the model is relatively accurate, with most of the data being fit. Variables with positive coefficients are more likely to be associated with a rating of five while variables while negative coefficients are associated with a non-five rating. As such, “fast”, “great”, “thank”, “good”, and “vendor” predict a five rating while “product”, “finalize”, and “early” are more likely to be associated with a non-five rating (see Table 7).




Figure 7: Error as a function of lambda

The odds ratios are interpreted as follows: a transaction which has “fast” in the feedback is 4.17 times more likely to be rated a five compared to a purchase review that does not. Therefore, the odds of a five rating are 130% and 115% greater when the words “great” and “thank” are found in the feedback than when they are not. In contrast, the presence of “product”, “finalize”, and “early” in buyer feedback increases the odds of a non-five rating by 2.1%, 3.1%, and 27.1%, respectively.

Table 7: Logistic lasso coefficients and odds ratios

Variable	Coefficients	Odds Ratios
Intercept	3.06542	21.4434
Fast	1.42956	4.17687
Great	0.83294	2.30007
Thank	0.76685	2.15298
Good	0.01742	1.01757
Vendor	0.00872	1.00876
Product	-0.0214	0.97886
Finalize	-0.0311	0.96937
Early	-0.3159	0.72917

Discussion

Based on the descriptive text analysis, “finalize” and “early” are the most frequently occurring words regardless of product type, rating, and price. This indicates that buyers put a fair amount of trust in vendors before the transaction is completed (i.e. receiving the product they have purchased). The full extent of early finalization will be later discussed. Nevertheless, based on the cluster dendrograms and word associations, buyers that have rated their transaction a five often reported the stealth and speed of the delivered product as reasons for their satisfaction. In contrast, buyers that did not rate their transaction a five maintained that the vendor had scammed them, reporting that the package had not arrived or was empty. However, the cluster dendrogram did also reveal that some non-five-rated transactions were “good”. More generally, however, feedback from these buyers were predominantly neutral, containing words such as “still”, “shipping”, “day”, “update”, and “order”.

This is corroborated by the results of the sentiment analysis. Indeed, while five-rated transactions possessed a positive average polarity, non-five-rated transactions were only marginally negative, verging on neutrality. To this extent, Abraxas buyers do not often harshly criticize vendors and products they are dissatisfied with. In contrast, the sentiment analysis reveals a clear difference between transactions that were finalized early that those that were not. To this extent, the finalize early corpus had a lower average polarity (0.31) than the non-finalize early corpus (0.66). This suggest that those who finalized early were less likely to use positive verbiage in their feedback relative to those who did not finalize early. Finally, the logistic lasso regression revealed that words such as “fast”, “great”, “thank”, “stealth”, and “good” predicted for a five rating. As such, the presence of these words in a feedback were associated with a five rating by a buyer. This demonstrates further congruity between a buyer’s word choice and the rating they rendered.

Finalize Early and Consumer Satisfaction

On cryptomarkets, there are two methods for fund exchange: 1) holding monies in escrow until product delivery and 2) finalizing early, forgoing escrow and transferring the funds upon purchase (Martin, 2013). As we can see on Abraxas, finalize early is often used, constituting 23.3% (1130) of all transactions in the dataset. The extensive use of finalize early tells us a great deal about trust on cryptomarkets. While finalizing early does not always equate to a five-rated transaction as 46 transactions that were finalized early were below five, it does reveal quite a bit about the nature of trust on Abraxas. To this extent, the decision to do business with a vendor on the part of buyers is governed not by a blind trust but an informed one. Indeed, buyers seemingly depend on vendor reputation when making the decision to finalize early. They are presumably using the information on a vendor’s previous transactions as a yardstick for trust, taking a risk by forgoing escrow based on the information available to them. This, moreover, reveals a great deal about information asymmetry on cryptomarkets.

Nevertheless, while finalizing early is seemingly correlated with trust, it is not necessarily correlated with consumer satisfaction. Indeed, it appears, on one hand, that the average rating between transactions that were finalized early and those that were not did not differ very much. On the other hand, the sentiment analysis revealed that buyers that did not finalize early were more likely to use positive verbiage relative to buyers that finalized early.

The suggestion here is that trust and consumer satisfaction are not necessarily correlated on Abraxas as while a buyer might trust a vendor enough to finalize their transaction they still may not be satisfied with the transactions once it is completed.

From this, it is evident that information asymmetry is not overcome on Abraxas. While buyers may subscribe to market-level information on who is trustworthy, ultimately finalizing their transaction early, this does not always equate to customer satisfaction. Indeed, as with licit markets, there is no guarantee of consumer satisfaction. Although the rating and feedback systems on Abraxas reduce information asymmetry a fair amount, this market, like any illicit market, is subject to the exigencies of uneven information flows. While a buyer can be confident that transacting with a reputable vendor will be a pleasant experience, there is no guarantee that the transaction will be as such.

This raises questions about why buyers might choose to finalize early in the first place. If a pleasant experience is not guaranteed, what factors would necessarily encourage a buyer to trust a vendor? Moreover, while it is understandable that a buyer who has done business with a vendor on previous occasions might finalize a transaction early, it is unclear why a buyer with no previous history with a vendor might choose to do so. Based on these findings, it is likely that buyers are choosing to rely on the shared experiences of other buyers who have engaged with a vendor, using the established reputation of vendor to make the decision to finalize early. As such, assumed reductions in information asymmetry pursuant to the public availability of buyer feedback might encourage prospective buyers to not only do business with a vendor but finalize early.

However, we cannot rule out social convention or pressure as possible explanations for why buyers might choose to finalize early when engaging with a vendor that they have little history with. In short, because many buyers on Abraxas choose to finalize early, this might encourage other buyers to do so. Finalizing early, then, might be viewed by prospective buyers as a transactional convention given its popularity. It may also be the case that the inherent trustworthiness of a vendor might endear them to current and future buyers who would engage in early finalization despite the risks associated with this action. Regardless, it is not possible to determine why buyers choose to finalize early without explicit data on this topic. Future qualitative research on cryptomarket buyers might endeavour to query buyers on their reasons for finalizing early.

Information Asymmetry on Cryptomarkets

From both the sentiment analysis and lexical predictors of five-rated transactions, it does appear that the sentiment structure of qualitative reviews differs between five-star and non-five-star ratings. However, this does not necessarily mean that buyers who did not rate their transaction a five felt negatively about the transaction. While zero and one-rated transactions were harshly criticized by buyers due to their disappointment with the product (i.e. the product did not arrive or was not as advertised), two, three, and four-rated transactions typically received value-neutral feedback. This suggests that consumer satisfaction on Abraxas is not subject to a particularly strenuous or hard-to-achieve standard. Buyers will overwhelmingly rate a transaction a five if the product arrives on time and is as

advertised by the vendor. To this extent, it was not surprising to see that words such as “fast”, “great”, “thank”, “stealth”, and “good” were associated with a five rating.

Many markets, criminal and licit, are often plagued by information asymmetry whereby knowledge about goods and services are not uniform between buyers and sellers. This can lead to a market for lemons where consumers possess less valid or reliable information about the quality of the goods relative to vendors. According to Herley and Florenio (2009), the uncertainty created by low-quality vendors imposes a tax on every transaction conducted in the market. That is, high-quality vendors stand to make less as the presence of low-quality vendors both discourages buyers from engaging in transactions and drives down the price of goods and services. This general uncertainty created by the presence of bad faith vendors imposes a tax on every transaction conducted in the market.

According to Thomaz et al. (2020), “cryptomarkets consists of two tiers of players: knowledgeable experts with more information, and newcomers who not only have less information but also do not know how to weigh sources of information and reputational cues properly”. This creates a fascinating situation where these naïve players abide by the rules established by knowledgeable players while the knowledgeable players put little to no trust in others. Moreover, these knowledgeable actors benefit from new market entrants who are otherwise naïve to the conditions of the market. A market will disintegrate in the absence of trust, but trust must be first predicated on the spread of information on who is trustworthy. In essence, trust is predicated on information. Indeed, buyers on Abraxas, whether naïve or competent, must accumulate enough evidence to convince themselves of the likelihood of a positive outcome: that they are likely to get what they paid for, and not get their money stolen or be arrested for their activities. As such, entire narratives are created around the trustworthiness of specific vendors.

To this extent, finalizing early and the Pollyanna principle are possible features of healthy information flows on Abraxas; a lack of information asymmetry. To this extent, the predominance of five-rated transactions is possibly due to the singling out of trustworthy and reputable vendors through a natural process of selection and elimination. As such, information on a vendor is created by past buyers who either applaud or chastise their respective transaction with said vendor. Indeed, a prospective buyer’s knowledge and desire to transact with a vendor is based on what prior buyers have suggested. We can think of this as an information cascade whereby buyers’ preference for a particular vendor is compounded across multiple transactions.

As demonstrated, this information cascade produces an arbitrary pareto distribution where a small number of reputable vendors account for a majority of transactions. Thus, the decision to finalize early by Abraxas buyers may stem from these information cascades where a trusted few are held in high regard. Once an information cascade is underway it is difficult to stop or reverse. Practically speaking, this means that reputable buyers will continue to remain reputable in the eyes of buyers while disreputable buyers will also remain as such. This description is corroborated by findings made by Janetos and Tilly (2017) as well as Batikasa and Kretschmer (2018). In short, both studies demonstrate that the length of a vendor’s transactional experience reduces the likelihood of market exit as a longer

transactional history is correlated with continued market participation. Once a vendor is marked early as trustworthy or untrustworthy, they will likely remain as such over their tenure on the market.

Functionally, these information cascades serve to minimize information asymmetry on Abraxas as buyers are made aware of which vendors are reputable. In contravention to information asymmetry, this level of certainty created by the presence of reputable vendors would not impose a tax on every transaction conducted on the market. Buyers and vendors are able to transact in a fairly transparent and collegial manner. For buyers, they are perhaps less likely to be afflicted by the fear of fraud or duplicity on the part of vendors. As such, in cryptomarkets like Abraxas, buyers are more likely to have a pleasant experience compared to terrestrial markets where less checks and balance that mitigate fraudulent vendor activity exist. This might perhaps explain the presence of the Pollyanna principle on Abraxas where the majority of transactions receive a 5-star rating. Indeed, buyers are simply reporting on their pleasant experiences on cryptomarkets relative to less pleasant experiences on terrestrial markets as opposed to inflating their scores given the abundance of 5-star ratings on the market.

Conclusion

As a theoretical matter, this study demonstrates the validity of collective or social learning on cryptomarkets. Indeed, the use of reputation systems on these platforms facilitates the spread of knowledge. This would ordinarily be difficult if first-hand experiences were not properly catalogued and made available to those on the markets. As such, prospective Abraxas buyers can observe the feedback of previous buyers and update their beliefs on the quality of a vendor. Of course, the value of a consumer feedback is often contingent on the volume and consistency of feedback. While several successful transactions may be a stroke of luck, a long history of successful transactions is a good indicator of reliable business practices and high-quality products.

This is the first such study to examine the lexical predictors of customer satisfaction and dissatisfaction on cryptomarkets using text mining. As such, future research should seek to replicate these analyses on other cryptomarkets to test the generalizability of these findings. There are, moreover, a bevy of methodological approaches that might be pursued in these future studies. These include exploring associations between buyer sentiment and vendor ratings using modified pre-processing techniques, applying sentiment analysis to product categories, and applying techniques such as random forest to predict emergent trends in vendor selection on the part of buyers.

Chapter 4: Testing the Efficacy of Six Simulated Targeting Strategies on a Dark Web Cryptomarket

Emerging threats from cyberspace have engendered proactive efforts from law enforcement to curtail and counteract these malicious actors and organizations. These strategies and tactics will differ in their operational parameters and functional objectives. Some interventions may seek to stem the flow of new actors while others aim to thoroughly dismantle the structure of a criminal organization in its entirety (Morselli, 2009). Regardless of the nuances of the intervention, one fundamental question remains: how effective was it? This question of what works and what does not has not been extensively applied to the study of cryptomarkets.

A relatively new criminal phenomenon (Martin, 2013a), cryptomarkets are illicit online marketplaces which facilitate the truck, barter, and trade of various and sundry illicit goods and services among buyers and vendors. While governments and law enforcement agencies worldwide have made numerous attempts at disrupting the ease of operation of these illicit entities they have metastasized, adopting new methods for both securitizing their continued operation and expanding their scope and influence (Shortis, Aldridge, and Barratt, 2020). As such, it is unclear whether these interventions have had a pronounced impact on the cryptomarket ecosystem. Nevertheless, a growing number of studies (Malm & Bichler, 2011; Natarajan, 2006; Wood, 2017) have provided empirical evidence on the utility of social network analysis in understanding the structural composition of criminal organizations. Moreover, cryptomarket scholars (Duxbury and Haynie, 2018; 2020) have begun testing the efficacy of strategic interventions.

While traditional methods of targeting criminal networks have prioritized the identification and removal of “kingpins”, Morselli (2009) contends that “the fluidity and flexibility of the structure of certain illicit networks makes them resilient to traditional law enforcement strategies”. Crucially, Duxbury and Haynie (2020) note that prior research (Holt, Strumsky, Smirnova, & Kilger, 2012; Kenney, 2007; Morselli, 2009; Malm & Bichler, 2011; Natarajan, 2006; Wood, 2017) which has used social network analysis to measure the structure and actors within a criminal network have failed to apply supplementary simulation methods that isolate probable vulnerabilities in the criminal network. In short, these studies have not accounted for probable network adaptation following intervention. This can be applied to generate informed strategies that are better able to disrupt the operation of criminal networks (Duxbury and Haynie, 2019). Thus, studies which leverage computer simulations to understand the impact of strategic interventions must consider probable adaptation on the part of actors within the network.

Utilizing sequential node deletion, this study examines the efficacy of six different targeting strategies (lead k, eccentricity, total revenue generated, cumulative reputation score, listing amount, and random targeting) in disrupting the ease of operation of a dark web cryptomarket. To this extent, five outcome variables (number of isolates, number of components, average number of nodes in components, average geodesic distance, and number of nodes in the largest component) are used to measure the performance of each targeting strategy. This study will set parameters to govern the purported behaviour of actors when nodes are removed. As such, the transactional network’s overall behaviour can be accurately modelled (Bright et. al, 2017) through an evidence-based calculus.

Literature Review

When dealing with criminals, law enforcement officials are constantly playing a game of cat and mouse where adaptations made by criminals force law enforcement to make counter-adaptations. This iterated game is particularly prominent within cyberspace where innovation and progress are the norm (Wall, 2001). Nevertheless, this should and does not stop law enforcement from testing new strategies for disrupting criminal activity and destabilizing criminal networks. Williams (2007) and Thraxter (2010) contend that law enforcement, while often slow in responding to emergent cyberthreats, have increased their efforts to curtail these burgeoning threats. This is no different for cryptomarkets. According to Martin (2014a), law enforcement organizations worldwide have undertaken a number of interventions to destabilize these illicit online marketplaces, including market infiltration and digital forensics, vendor arrests, mail scanning, and market takedown.

Of course, the principal dilemma facing law enforcement agencies tasked with combating cryptomarkets is where to target scarce resources. This is further complicated by the exigencies of the highly volatile cryptomarket environment. Indeed, the short lifecycle of these illicit marketplaces (Christin, 2013; Christin, 2015) makes the investigation of these entities particularly difficult as a market may cease to exist before the investigation is completed. Furthermore, constantly improving security and encryption protocols compounds the difficulty of adequately policing these entities. According to van Buskirk et al. (2014), “administrators are heeding the lessons of prior market closures and are taking extra steps to fortify their sites against external penetration” (54). In short, the task of disrupting the ease of operation of cryptomarkets grows increasingly difficult while the resources required for such undertakings remain scarce.

However, criminologists (McGloin & Rowan, 2015; Shaw & McCay, 1942; Warr, 2002) have increasingly observed that a large portion of criminal activity is group-based. Indeed, criminals often do not act alone but are instead imbedded in a network of similarly motivated actors. To this extent, researchers have increasingly relied on social network theory and methods to understand the structure, operation, and vulnerabilities of criminal entities of varying size (Kennedy, 2008; McGloin, 2005; Papachristos, 2009, 2011, 2014). Indeed, a myopic focus on single individuals on the part of law enforcement is unlikely to bear fruit when combatting an association of actors. To this extent, a growing area of research examines how criminal networks respond to disruption (Duijn et al., 2014; Malm and Bichler, 2011; Morselli, 2009).

Consider research by Krebs (2002) which documented the structural properties of the 9/11 terrorist attack. According to the author, the criminal network’s extreme diffusion made it particularly resilient to disruption. Indeed, any single actor could only incriminate a maximum of four other members of the network if identified and arrested. In contrast, Wood (2017), examining the structure of an international heroin trafficking network, found that the removal of 20% of all actors had a considerable disruptive effect on the network.

Topological Features of Criminal Networks

For the resilience and behaviour of a criminal network to be understood, researchers have turned to examining a network’s topology (Duxbury and Haynie, 2018). Indeed, the structure of network will often determine how it responds to law enforcement intervention. Importantly, the topology of a criminal network is often unique and is organized based on differences in security concerns and constraints on the efficient mobilization of resources (Raab and Milward, 2003). No two criminal networks are the same. According to research from Morselli et al (2007), drug distribution networks typically rely on a hierarchical network

structure where high-profile distributors insulate themselves from the brunt of the network activity by connecting to only a few actors. This ensures a certain level of protection as constant exposure to other actors and elements within the network increases the likelihood of arrest.

This is not the same for all networks. Alternatively, research by Diekmann et al. (2014) and Stephen and Toubia (2009) on social commerce networks revealed that these networks are premised on preferential attachment. As such, these networks possess scale-free properties where a small number of desirable, trusted actors retain a large proportion of customers within the market. Such power law dynamics are not particularly reliable in the criminal context as scale-free criminal networks are more vulnerable to crippling targeted interventions given the presence of highly connected vertices (Albert et al., 2004). In fact, research (Raab and Milward, 2003) suggests that criminal networks will often veer away from highly centralized topologies to ensure greater structural robustness when confronted with interventions against them.

However, this aversion to scale-free properties is not universal across all criminal networks as some network topologies will naturally abide by power law dynamics. Indeed, preferential attachment is unavoidable in environments where trust is scarce and difficult to establish. Past research demonstrates that skewed degree distributions are also a characteristic feature of criminal networks. Among cannabis cultivators, Duijn et al. (2014) demonstrate a pronounced power law distribution where a small number of actors produced and traded a disproportionate amount of cannabis. In a similar study of a drug trafficking network, Natarajan (2006) discovered a small number of disproportionately high degree traffickers among a large contingent of actors with low degrees. This is similar to Varese (2010) whose examination of a Russia Mafia group in Italy revealed a heavy-tailed degree distribution. As such, Varese concluded that the group was hierarchically structured and polycentric. While Krebs (2002) demonstrated that the 9/11 terror cell was a diffuse network, this is not the case for all terrorism networks. The degree distribution can vary considerably in a terrorism network. For example, Morselli et al. (2007) and Qin et al. (2005) found truncated power law distributions in the terrorism networks they examined.

According to Newman (2002), the removal of a highly connected vertex often fractures the network into numerous distinct components. This is reflective of degree-mixing patterns which characterize scale-free networks (Newman, 2003). To elaborate, Alm and Mack (2017) contend that networks where high- k actors are connected to low- k actors (degree mixing or disassortative mixing) are more susceptible to the disruptive impact of key vertex removal relative to networks where degree mixing does not occur (assortative mixing). Wood (2017) documented disassortative mixing in several drug distribution networks, corroborating Kennedy's (2008) contention that many real-world drug markets are highly susceptible to law enforcement intervention. As it relates to cryptomarkets, Barratt et al. (2016a; 2016b), originally speculated that these online markets were subject to low disassortative mixing as many buyers reported experimenting with new products and vendors. Moreover, there is relatively less risk in purchasing goods and services through online markets.

Simulated Interventions on Criminal Networks

Given the inherent difficulties associated with the procurement and cleanliness of criminal network datasets, studies on simulated law enforcement interventions are scarce relative to studies documenting the structure of criminal networks. Regardless, much can be gleaned from the studies that have evaluated simulated interventions against criminal

networks. One such study by Keegan et al. (2010) contrasted the structural robustness of a drug trafficking network with a proxy gaming network. Applying k-based sequential node removal to each network, the authors observed that removing the top 5% of nodes based on degree centrality dismantled both networks whereas the random removal of 5% of nodes failed to yield a comparable result. In short, Keegan et al. (2010) demonstrated the disruptive impact of k-degree removals in both licit and illicit networks. In contrast, Xu and Chen (2003) used a simulation methodology to examine terrorist, methamphetamine trafficking, and gang networks. The authors concluded that the targeting of hubs and brokers was ineffective in disrupting the network's ease of operation. In contrast, strategies which prioritized the targeting of brokers proved more effective as these specific actors were responsible for keeping the network together. Together, these studies demonstrate that differences in the topology of a criminal network will yield different vulnerabilities which require different strategic interventions. Indeed, one targeting strategy will not be equally effective on all networks whose structural compositions differ.

Applying computer simulations to evaluate the impact of law enforcement interventions on two drug trafficking networks, Bright, Greenhill, and Levenkova (2010) focused on the removal of hubs. As with previous studies, the researchers found that the removal of key actors by law enforcement can create a relatively speedy structural collapse. Furthermore, Bright, Greenhill, Britz, Ritter, and Morselli (2017), investigating the effectiveness of six law enforcement intervention strategies against a drug market against three outcome measures, found that the removal of actors with betweenness centrality was the most effective strategy. This was followed by removing actors who made the most money.

Examining four criminal networks with varying network structures, Duxbury and Haynie (2019) applied agent-based modelling to evaluate how criminal networks respond to disruptions. The authors made two important conclusions. First, "isolated law enforcement disruptions maybe unsuccessful at reducing future levels of crime in efficiency-oriented networks" (Duxbury and Haynie, 2019, 335). Second, disruption tends to yield time-persistent damage to a network which prioritized security. This suggests that future law enforcement interventions should attack security-oriented criminal networks.

As it pertains to cryptomarkets, Duxbury and Haynie (2018), building off their prior research (Duxbury and Haynie, 2017) on transactional networks on the dark web, conducted disruption simulations on an opioid market. Their results demonstrated that the removal of high centrality actors in repeat transactions yielded a decrease in the size of the largest network components. However, isolated groups and potential components increased in size. According to the authors (Durbury and Haynie, 2018), "these results suggest that targeting any available combination of high-profile distributors may be an alternative strategy to leading distributor removal when leading distributors are difficult to isolate or identify" (245). Consistent with research in drug distribution networks (Carley, 1995; Duijn et al., 2014; Morselli et al., 2009; Wood, 2017), the authors found that removing the most prolific vendors in sequential order fragmented the network in relatively little time.

The same authors (Duxbury and Haynie, 2020) conducted a second study which applied computer simulations to test the responsiveness of a dark web drug network. The researchers found that "while targeted attacks were effective when conducted at a large-scale, weak link and signal attacks deter more potential drug transactions and buyers when only a small portion of the network is attacked" (34). They also found that intentional attacks were

generally more effective as actors grow more cautious about forging ties when the network is attacked. Under these conditions, network robustness is undermined in the long term.

Research Questions

Building off research from Duxbury and Haynie (2018; 2020), this paper seeks to answer three research questions:

1. Of the six proposed disruption strategies, which offers the greatest initial amount of damage to the Abraxas transactional network?
2. Of the first 100 nodes that are removed per each disruption strategy, does their impact carry-over across all outcome measures?
3. What do these strategies tell us about the efficacy of dark web disruption strategies?

Given the dearth of research on this topic, scholars and law enforcement are generally uncertain about the measurable impact of cryptomarket disruption strategies. The effectiveness of cryptomarket intervention strategies is an area where knowledge is lacking (Shortis, Aldridge, and Barratt, 2020). While Duxbury and Haynie (2020) have applied sequential node removal to one cryptomarket, it is unclear how generalizable these findings are. As criminal networks are adaptive and dynamic, different disruption strategies are likely to yield different results. When it comes to cryptomarkets, it is unclear which strategies work, and which ones do not. The first and second research questions will address this important gap in the scholarly literature, comparing and contrasting the effectiveness of six different disruption strategies across five impact measurements. As such, the explicit focus of this chapter is not on the structural robustness of cryptomarket transactional networks, but on the efficacy of strategic interventions which might be tried against these networks.

This expands on Duxbury and Haynie's (2020) study which leveraged three intervention strategies (high k vendors, low k buyers, and vendor rating) across three impact measurements (number of ties, numbers of transactions, and network density). To this extent, the objective is to determine which strategies are most effective across single outcome measures and across all measures. In addition, this represents a novel opportunity to identify inherent differences in each strategic intervention. Indeed, it may be the case that while each strategy possesses a different targeting objective, they may target the same actors within the network. As a result, these interventions, while purported to be strategically distinct, are functionally similarly if not identical.

The third research question seeks to leverage the findings of the first and second research questions to speculate on the overall efficacy of law enforcement interventions. The structure of a criminal network naturally lends itself to the generation of disruption strategies. In short, this question strives to evaluate how different criminal network disruption strategies might affect the immediate and long-term impact of dark web criminal networks. Network activity in the aftermath of disruption provides insight into how criminal organizations behave in unstable environments. Social network theory and analysis is ideally suited to understanding how disruption efforts affect crime groups' behaviour, coordination, and time of recovery. Moreover, the combined use of social network analysis and computer simulations overcome the well-known methodological and data collection problems associated with examining dark networks (Bright and Delaney, 2013; Bright, Koskinen, & Malm, 2018; Morselli, 2009; Wood, 2017). Thus, in addition to making theoretical advancements in understanding organized crime and informing criminal intelligence, this question will also provide methodological contributions, demonstrating the utility of

computational methods and social network analysis in understanding criminological phenomena.

Data

As with previous chapters, I again use a buyer-seller dataset from the Abraxas cryptomarket (Branwen et al., 2015). Apart from the anonymous cryptomarket analysed by Duxbury and Haynie (2017; 2019), this is the only marketplace where unique identifiers are available for buyers. From the 5434 illicit transactions, a single two-mode network featuring vendors and buyers was created. Vendors were identified based on their unique vendor name while buyers were identified based on their HTML code. As such, the transactional network consisted of 5434 pairs, with 269 unique vendors and 2794 unique buyers. This analysis used directed ties.

Table 1 present the descriptive network statistics of Abraxas' transactional network. First, the network is diffuse with a network density of 0.0007. As such, only 0.07% of all possible transactions occurred. Furthermore, the full network consists of 29 components, with one component containing 97.6% (2726) of all nodes within the network. The remaining connected components consisted of 19 dyads, 7 triads, and single assortments of components of various sizes. As expected, there are no isolates as a transaction must involve both a buyer and a vendor. Nodes within the Abraxas transactional network, based on the eccentricity measurement, have a maximum distance of 11.23 from one another, on average. Comparable mean values can also be observed for vendors (10.32) and buyers (11.33).

Table 1: Network characteristics

Network Characteristics	Mean (SD) or Total	Range
Unique Actors/Nodes	2794	-
Unique Vendors	269	-
Unique Buyers	2525	-
Isolates	0	-
Total Unique Edges	3935	-
Density	0.0007	-
Indegree	2.15 (2.2)	1-34
Outdegree	20.2 (39)	1-330
Indegree Centralization	0.01	-
Outdegree Centralization	0.12	-
Eccentricity (All)	11.23 (1.9)	1-16
Eccentricity (Vendors)	10.32 (3.38)	1-15
Eccentricity (Buyers)	11.33 (1.64)	1-16

Methods

By virtue of their orientation in a network, the behaviour of each individual node is dependent on the behaviour of every other node within the network (Bright et al., 2017). Simply because two nodes are unconnected does not necessarily mean that they do not affect one another in some capacity. Indeed, downstream effects are plausible (Newman, 2003) as the removal or inclusion of new nodes changes the dynamic of a network, and by extension, the behaviour of the actors within it (Namatame and Chen, 2016). All told, criminal networks

are comprised of individual actors whose mutual relationships strengthen or dissipate when internal and external stimuli are added to the network. Thus, criminal networks are not static entities. The structure of a network may change based on the behaviour of the actors within it. This fact cannot be ignored when examining the efficacy of interventions on criminal networks.

Simulation of Interventions Against Criminal Networks

Research into the disruption of criminal networks has posited several methods of reducing the ease of operation of these networks. In curbing the network activity of a drug market, Kennedy (2008) advocates for a “focus deterrence” strategy which simultaneously removes multiple influential criminals in a single stroke. This is done to reduce the likelihood of a resultant power vacuum by removing the actors most responsible for activity within a network. This method of targeting the most influential actors is challenged by methods which maintain that network disruption is best achieved by targeting brokers within the network (Burt, 2000). These brokers bridge structural gaps in a network, connecting segments of a network via the maintenance of pathways. This is particularly important for strategies targeting gangs as greater disruptive impact may be achieved if actors spanning local network clusters are targeted for removal. Indeed, this strategy posits the targeting of connectors as opposed to distributors.

Based on numerous studies examining real and hypothetical law enforcement interventions (Morselli, 2009; Wood, 2017; Alm and Mack, 2017; Duxbury and Haynie, 2020), the most common method of testing the structural robustness of a network is to “remove vertices in descending order of magnitude and to measure the proportion of network features as a function of the actor’s removal” (Duxbury and Haynie, 2018, 245). Importantly, this paper does not endeavour to measure the structural robustness of Abraxas, but rather the efficacy of proposed strategic interventions which target the actors therein. Nevertheless, this method can also be applied to measure the efficacy of strategic interventions.

Furthermore, this study deviates from other studies (Haynie, 2018; 2020) measuring cryptomarket intervention as it places each targeting strategy into one of two categories for disrupting criminal network. Based on previous research (Bright et al., 2017), strategies for criminal network disruption can be divided into two categories: the network approach and the human capital approach. The network approach focuses on individual actors that occupy strategic positions within criminal networks (Sparrow, 1991; Klerks, 2001; Schwartz and Rouselle, 2009). These predominantly revolve around common centrality measurements such as degree centrality, betweenness centrality, and eccentricity. Originating in economics, human capital encompasses “the competencies, knowledge, social and personality attributes, including creativity, embodied in the ability to perform labour so as to produce economic value” (Duijin, Kashirin, and Sloot, 2014, 4236). As it pertains to illegal markets, human capital is assembled and integrated in the form of trust.

I employ sequential node deletion pursuant to six law enforcement strategies: lead k (degree centrality), eccentricity, unique items bought/sold, cumulative reputation score, total purchase price, and random targeting. Mirroring Bright et. al (2017), each strategy is premised on the hypothesized aims of law enforcement agencies and fall under either a social or human capital approach. Each targeting strategy begins with the full 2794 actors within the network then deletes one node at a time based on the strategic objective of the intervention. Isolates are then given the choice to rejoin the network before calculating the output

measures. This simulation strategy was selected due to its successful use by Bright et al. (2017) and Gilbert & Troitzsch (2005). The six targeting strategies are as follows:

1. Random targeting: targets are selected at random regardless of their role in the market. This targeting strategy possesses no overt strategic objective. It is premised on opportunistic intervention.

Interventions that Target Only Network Capital

2. Degree centrality targeting: lead k actors are removed in descending order. Within Abraxas, these are the actors with the highest number of trade partners. This is a fairly standard measurement by which network-based node removal is conducted.
3. Eccentricity targeting: nodes in the network will be removed based on their distance from a specific node to any other node. Eccentricity measures the maximum distance of one node to any other node in the network. As such, the eccentricity of a node in a connected network is the maximum distance between that node and another over all nodes in the network.

It is important to note that betweenness centrality, while a staple of the network capital approach, was not featured as a targeting strategy as the Abraxas transactional network did not contain influential brokers which connected disparate parts of the network. In general, the directness of cryptomarket transactions does not allow for the existence of brokers as would be present in terrestrial criminal markets.

Interventions that Target Only Human Capital:

4. Unique items bought/sold targeting: nodes are removed based on the number of unique items bought or sold by an actor.
5. Total purchase price targeting: nodes are removed based on the total revenue generated or spent by an actor.
6. Reputation targeting: deletions are based on the cumulative reputation score of actors.

Finally, I use five outcome variables to assess the efficacy of each strategy:

- 1) Mean geodesic distance in the network
- 2) Number of nodes in the largest components in the networks
- 3) Average number of nodes in components
- 4) Number of components
- 5) Number of isolates.

The first outcome variable examines the mean of the shortest path lengths between any two actors in the network. Smaller mean geodesic distances indicate that information and resources can travel more quickly throughout the network, promoting criminal activity. Thus, increases in mean geodesic distances indicate greater network damage. The second, third, and fourth variables measure network hierarchy and actors' integration into a centralized organization. Thus, they provide a measure of hierarchical network cohesion, where decreases in the size of the largest component, the average number of nodes in components, and the number of components reflect greater network damage. The fifth variable measures the fragmentation of the market based on the number of nodes without a tie. As the number of isolates increases, the network grows more fragmented and is generally less capable of achieving organizational goals as a cohesive unit.

It is, nevertheless, important to clarify differences between this simulation strategy and those pursued in other studies measuring the impact of interventions on cryptomarkets. Duxbury and Heynie (2019) leveraged three intervention strategies in their first study of a dark web opioid network: 1) high k vendors, 2) low k buyers, and 3) vendor rating. The impact of these interventions was measured across three impact measurements: number of ties, numbers of transactions, and network density. In their second study, Duxbury and Haynie (2020) used three attack strategies: 1) weak link attacks that delete large numbers of weakly connected vertices, 2) signal attacks that saturate the network with noisy signals, and 3) targeted attacks that delete structurally integral vertices. These interventions were measured across the number of ties, network density, and number of isolates within the network.

While this study shares some targeting strategies (lead k and vendor reputation) and outcome measurements (number of isolates) with the aforementioned studies, it provides a wider array of targeting strategies and outcome measurements that have not been attempted. As such, this study offers a more in depth look at the efficacy of cryptomarket targeting strategies, building upon prior research on this topic by more closely examining the relative and comparative impact of each targeting strategy. This adaptive simulation strategy qualifies as the most extensive within the scholarly literature on cryptomarkets. Moreover, the use of network and human capital frameworks adds a more rounded analytical focus, segmenting the targeting strategies based on a higher order functional objective premised on network position or human competency. This has not been attempted in prior cryptomarket simulation studies.

Like Bright et al. (2017) and Gilbert & Troitzsch (2005), I perform 100 iterations of the simulation for each target strategy. Each outcome measure is then averaged over the 100 runs to produce plots of the average value over time (Berk, 2008; Birks & Davies, 2017; Birks et al., 2012; Groff et al., 2018; Weisburd et al., 2017).

Accounting for Network Adaptation

Real-world data on criminal networks are typically drawn from captured networks, rendering observations of the network before and after disruption almost impossible (Bright et al., 2018; Morselli, 2009). As such, sequential node deletion simulations must incorporate network adaption and preferential selection processes that are premised on some sort of ground truth.

Following Bright et. al's (2017) adaptation procedure, "network adaptation was modelled by giving the network an opportunity to replace an actor that was removed due to sequential node deletion". In this study, I assumed that replacement actors should possess three necessary characteristics: 1) the same product bought/sold as the deleted actor, 2) the same shipped to/from location as the deleted, and 3) the highest possible reputation score of all eligible replacements. Each of these replacement criteria were weighed the same. In other words, replacement actors must match the base-level profile of the deleted actor while also possessing a relative high level of trustworthiness such that surrounding nodes would comfortably do business with them. Once a node had been removed in each sequential deletion, the first step was to identify how many nodes were made an isolate as a result of the deletion. Second, a single replacement node in the network which possessed the three aforementioned replacement characteristics was identified. Isolates were then given the opportunity to reconnect to the network via the identified replacement node. Importantly, the probability of reconnection was set to 0.5, indicating that the isolate had a 50% probability of

reconnection. All told, the network would replace an actor that was removed with the most suitable candidate. If a suitable candidate did not exist, the isolate did not re-join the network.

This adaptation process is based on network redundancy. Redundancy, in this case, refers to the number of different relationships between actors in a network. Importantly, the more redundancy there exists in a network the more viable options there are for replacing lost human capital. In short, replacements in Abraxas with a reliable reputation and suitable shipping country and product listing will serve as replacements once similar actors have been deleted from the market.

Findings

Simulation Results

Table 2 displays the results of all six simulated interventions across the five outcome measures. To facilitate comparisons across law enforcement strategies, I plot the five outcome measures on five separate graphs: number of active components (Figure 1(a)), number of isolates (Figure 2(b)), average number of nodes in components (Figure 1(c)), number of nodes in largest component (Figure 1(d)), and average geodesic distance (Figure 1(e)). All five plots show the results of simulations in which network adaptation is included. For each plot, the x-axis shows the number of steps performed, operationalized as the number of nodes deleted sequentially. At each step, one actor is removed. The y-axis reflects the specific outcome measure featured in the simulation.

Table 2 demonstrates the impact of deleting a single node per each intervention across all outcome measurements. Based on the number of isolates and components, it is readily apparent that eccentricity and random targeting are the least effective targeting strategies, producing the lowest average results per deleted node. Interestingly, degree centrality, reputation, total purchasing price, and unique items bought/sold each performed similarly across these two measurements. While the average is particularly stable for each targeting strategy across the average number of nodes in components, the number of nodes in largest component, and the average geodesic distance, clear differences are apparent based on the standard deviation and range. Again, eccentricity and random targeting are the least effective at disrupting the transactional network. Furthermore, degree centrality, reputation, total purchasing price, and unique items bought each perform similarly across these three measures, offering the greatest disruption per node deleted. Nevertheless, a closer look at speed of disruption for each targeting strategy across the five measures is warranted.

Table 2: Impact of Single Node Deletions by Strategy and Outcome

Measures	Initial Value	Mean	SD	Range
<i>Isolates^a</i>				
Degree Centrality ^b	0	1.77	4.46	0-91
Eccentricity ^b	0	0.03	0.18	0-3
Random ^b	0	0.03	0.18	0-3
Reputation ^b	0	1.77	4.46	0-91
Total Purchasing Price ^b	0	1.7	4.52	0-91
Unique Items Bought ^b	0	1.67	4.52	0-90

Components^a

Degree Centrality ^b	29	1.76	4.58	0-91
Eccentricity ^b	29	0.02	0.16	0-3
Random ^b	29	0.18	0.45	0-5
Reputation ^b	29	1.75	4.58	0-91
Total Purchasing Price ^b	29	1.69	4.64	0-91
Unique Items Bought ^b	29	1.66	4.63	0-90

Average Number of Nodes in Components^a

Degree Centrality ^b	96.35	0.04	1.35	0-70.48
Eccentricity ^b	96.35	0.04	0.06	0-2.72
Random ^b	96.35	0.04	0.16	0-3.12
Reputation ^b	96.35	0.04	1.35	0-70.48
Total Purchasing Price ^b	96.35	0.04	1.35	0-70.48
Unique Items Bought ^b	96.35	0.04	1.36	0-71.41

Average Geodesic Distance^a

Degree Centrality ^b	64.62	0.04	0.63	0-30.45
Eccentricity ^b	64.62	0.05	0.38	0-12.44
Random ^b	64.62	0.2	1.56	0-43.24
Reputation ^b	64.62	0.04	0.63	0-30.45
Total Purchasing Price ^b	64.62	0.04	0.51	0-23.24
Unique Items Bought ^b	64.62	0.04	0.51	0-23.24

Number of Nodes in Largest Component^a

Degree Centrality ^b	2726	0.98	6.47	0-169
Eccentricity ^b	2726	0.98	0.59	0-27
Random ^b	2726	0.98	1.31	0-39
Reputation ^b	2726	0.98	6.34	0-159
Total Purchasing Price ^b	2726	0.98	5.65	0-102
Unique Items Bought ^b	2726	0.98	6.11	0-117

^a indicates outcome measure; ^b indicates targeting strategy

As it pertains to the number of components, the maximal effect is measured as the highest number of components that are created upon intervention. In short, if the intervention is to be successful node deletion should yield a sizable increase in the number of components within the network (see Figure 1(a)). An increase in the number of components reflects network fragmentation and disruption of information flows. Upon closer examination, it is evident that degree centrality targeting yielded the fastest speed (by the narrowest margins) of relative disruption as the deletion of 251 nodes (9% of all nodes) yielded 2310 total components, a 7866% increase from the original 29 components. In comparison, reputation targeting yielded 2312 components after 244 nodes were deleted while total purchasing price targeting and unique items bought/sold resulted in 1948 and 2061 components after 288 and 422 nodes were deleted, respectively. Perhaps unsurprisingly, random targeting yielded a high of 269 components once 2525 nodes were deleted. Curiously, eccentricity offered the least disruption with a high of 38 components once 2702 (96.7%) nodes were deleted.

With regard to the number of isolates, the maximal effect is measured as the highest number of isolates that are created upon intervention (see Figure 1(b)). Reputation targeting appears to be the most effective strategy as the deletion of the top 299 nodes (10.7%) yielded 2202 isolates within the network. Degree centrality targeting offered the second fastest disruption with the deletion of the top 300 nodes yielding 2201 isolates. In this case, both strategies offered near identical results. Total purchasing price targeting and unique items bought/sold targeting created the third and fourth fastest disruptions, respectively. Total purchasing price yielded a high of 1841 isolates after 288 nodes were deleted while unique items bought/sold yielded 1559 isolates after 422 were removed. Eccentricity and random targeting offered the same maximal disruption, with 19 isolates created after 2765 nodes were deleted. These are particularly poor showings.

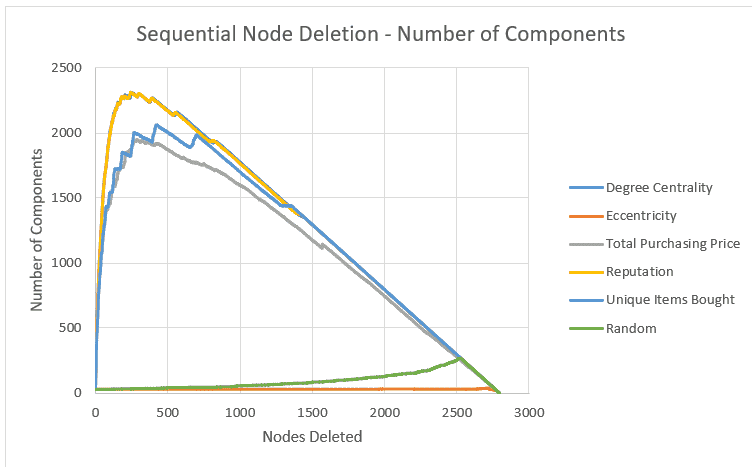


Figure 1(a): Number of Components




Figure 1(b): Number of Isolates

The average number of nodes in a component reflects the average size of components within the network. As such, maximal disruption is premised around reductions in the average number of nodes in components. The smaller the components the more fragmented the transactional network. Based on rank-order vertex removal simulations, the size of components plummets as the number of nodes is removed. This is particularly the case for degree centrality targeting, total purchasing price targeting, reputation targeting, and unique items bought/sold targeting follow similar pattern (see Figure 1(c)). Degree centrality targeting yielded the fastest disruption, with 44 (1.6%) node deletions reducing the average component size to 1.99 nodes (a 97.9% reduction from the original 96.35 average). Similarly, reputation targeting yielded nearly identical disruption as the rank-ordered deletion of 47 nodes reduced the average number of nodes within components to under two. In order to reduce the average number of nodes to 1.99, 70 and 71 nodes needed to be deleted for total purchasing price and unique items bought/sold, respectively. Random and eccentricity targeting yielded the slowest disruption as it required the random deletion of 2381 (85.2%) and eccentricity-based deletion of 2938 (98%) nodes to reduce the average component size to 1.99 nodes.

As it pertains to the number of nodes in largest component, the maximal disruptive effect is measured as the lowest number of nodes in the largest component following intervention. In other words, the smaller the largest component the more fragmented the network (see Figure 1(d)). Degree centrality targeting yielded the fastest speed of relative

disruption as the deletion of 562 nodes reduced the largest component to 3 nodes. In comparison, reputation targeting yielded the same result after 815 nodes were deleted. Total purchasing price targeting and unique items bought produced the same measure result (3 nodes) after 2621 and 1351 nodes were deleted, respectively. Random and eccentricity targeting yielded the same outcome after 2507 and 2736 nodes were deleted, respectively.

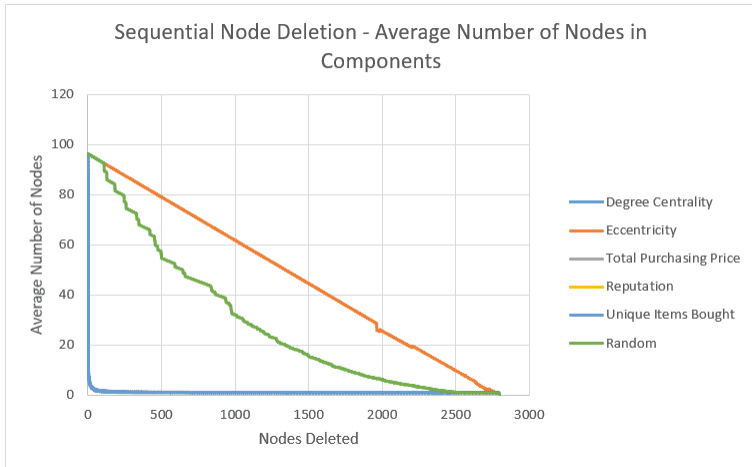


Figure 1(c): Average Number of Nodes in Components




Figure 1(d): Number of Nodes in Largest Component

While decreases in the average geodesic distance typically reflects improved communication between nodes as the distance from one node to all other nodes is short, this is not necessarily the case for the Abraxas transactional network. In this case, consistent decreases in the average geodesic distance pursuant to sequential node deletion are a product of a shrinking share of nodes which can be connected to. In short, the average geodesic distance decreases as there are less nodes to connect to. Reputation targeting appears to be the most effective strategy as the deletion of the top 99 nodes (10.7%) yielded an average geodesic distance of 0.41. Degree centrality targeting offered the second fastest disruption with the deletion of the top 96 nodes yielding an average geodesic distance of 0.93. Unique items bought/sold targeting and total purchasing price targeting created the third and fourth fastest disruptions, with the deletion of 226 and 230 nodes yielding an average geodesic distance of 0.8, respectively. Eccentricity and random targeting were again the least effective strategies, yielding averages of 0.5 and 0.3 once 2792 and 2240 nodes were deleted, respectively.




Figure 1(e): Average Geodesic Distance

Node Deletion Impact

Not every deleted node will have the same disruptive effect on a criminal network. By virtue of their influence and place within the network structure, the removal of specific nodes will have disproportionate impact on a network's ease of operation. Figures 2(a), 2(b), 2(c), 2(d), and 2(e) illustrate the percentage of node deletion by the percentage of disruption impact for each strategic intervention across all outcome measures. It is immediately apparent from these figures that the impact of node deletion can either be linear, curvilinear, or power law. A linear relationship means that node deletion and disruption impact are proportional, implying that the removal of a large number of nodes will result in an equally large level of disruption. A curvilinear relationship implies that a moderate number of deleted nodes account for a large amount of disruption impact. Finally, a power law curve implies that small number of deleted nodes accounts for an outsized portion of disruption impact.

Based on Figure 2(a), degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting are each based on a power law when it comes to the number of nodes in the largest component. In degree centrality targeting, 1% of deleted nodes accounted for 51.5% of disruption impact. This is also the case for the other three strategies as 1% of deleted nodes accounted for 46.2%, 50.8%, and 52.8% of disruption impact for total purchase price targeting, reputation targeting, and unique items bought/sold targeting, respectively. Random targeting is curvilinear with 10% of deleted nodes accounting for 26.6% of disruption impact while eccentricity targeting is linear. As it relates to the number of isolates (see Figure 2(b)), 5% of deleted nodes accounted for 45.2%, 45.8%, 45.2%, and 45.8% of disruptive impact for degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting, respectively. As such, each targeting strategy is governed by a power curve. Random and eccentricity targeting appear to be linear.




Figure 2(a)




Figure 2(b)

As it relates to the number of components (see Figure 2(c)), 5% of deleted nodes accounted for 46.3%, 47.3%, 46.9%, and 47.2% of disruptive impact for degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting, respectively. Each of these targeting strategies is governed by a power curve. Random and eccentricity targeting appear to be linear. Power curves are also present among

several targeting strategies as it relates to the average number of nodes in components (see Figure 2(d)). 1% of deleted nodes accounted for 86.3%, 85.8%, 86.2%, and 86% of disruptive impact for degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting, respectively. Interestingly, random targeting abides by a less pronounced power curve, with 1% of deleted nodes accounting for 41.3% of disruption impact. Unsurprisingly, eccentricity targeting follows a linear curve. Curiously, all targeting strategies abide by a power curve when it comes to measured impact on average geodesic distance. 1% of deleted nodes accounted for 73.8%, 85.8%, 74.6%, 60.3%, and 60.3% of disruptive impact for degree centrality targeting, reputation targeting, total purchase price targeting, unique items bought/sold targeting, and random targeting, respectively. 5% of the deleted node targeted based on eccentricity accounted for 46.1% of impact disruption.

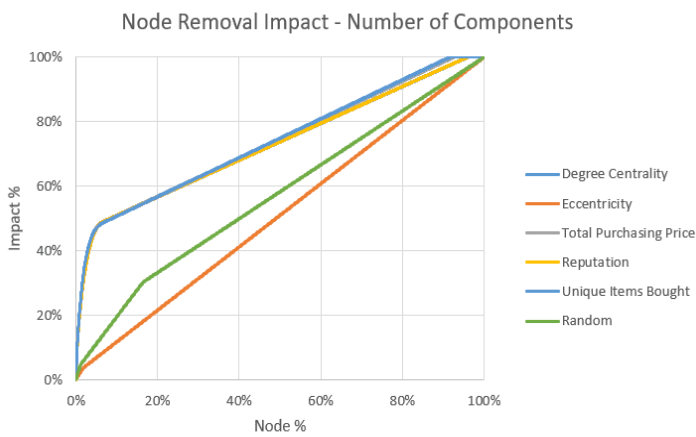


Figure 2(c)




Figure 2(d)




Figure 2(e)

Outcome Measure Carry-Over and Node Characteristics

Table 3 shows the percentage of the top 100 deleted nodes for each target strategy that are also held in common with each other respective targeting strategy for each outcome measure. In short, this table shows how many nodes within the top 100 simulated deletions are held in common by all targeting strategies. Based on these results, it is apparent that the

majority (> 50%) of deleted nodes are held in common among degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting for nearly all outcome measures. In fact, there appears to be a congruence of 90% or greater for isolates and components. This in part explains why the aforementioned disruption impact was so similar among these targeting strategies as the deleted nodes were the same actors. Moreover, these four targeting strategies did not share the same nodes with both eccentricity and random targeting. This again explains the sharp differences in their disruption performances as the same actors were not targeted for deletion as they were for degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting.

Table 3: Top 100 Actors Held in Common Across Targeting Strategies

Metrics	Degree Centrality	Eccentricity	Total Purchasing Price	Reputation	Unique Items Bought
<i>Isolates</i>					
Eccentricity	10%	-	-	-	-
Total Purchasing Price	95%	9%	-	-	-
Reputation	99%	10%	95%	-	-
Unique Items Bought	94%	10%	95%	95%	-
Random	7%	5%	7%	7%	8%
<i>Components</i>					
Eccentricity	6%	-	-	-	-
Total Purchasing Price	92%	5%	-	-	-
Reputation	100%	6%	92%	-	-
Unique Items Bought	93%	6%	95%	93%	-
Random	0%	3%	0%	0%	0%
<i>Average Number of Nodes in Components</i>					
Eccentricity	4%	-	-	-	-
Total Purchasing Price	85%	5%	-	-	-
Reputation	92%	5%	85%	-	-
Unique Items Bought	76%	4%	68%	73%	-
Random	2%	2%	3%	2%	6%
<i>Average Geodesic Distance</i>					
Eccentricity	0%	-	-	-	-
Total Purchasing Price	73%	0%	-	-	-
Reputation	86%	0%	74%	-	-
Unique Items Bought	49%	0%	45%	46%	-

Random	0%	7%	2%	1%	1%
<i>Number of Nodes in Largest Component</i>					
Eccentricity	2%	-	-	-	-
Total Purchasing Price	85%	2%	-	-	-
Reputation	97%	2%	85%	-	-
Unique Items Bought	78%	3%	78%	79%	-
Random	0%	2%	2%	1%	3%

Table 4 presents a complete array of descriptive statistics for the top 100 deleted nodes that are held in common across all outcome measures. Of the top 100 nodes deleted based on degree centrality targeting, 78 were held in common across the five outcome measures. Notably, eccentricity targeting and random targeted yielded no common deleted nodes across the outcome measures. Of the targeting strategies with commonly held deleted nodes, an overwhelming majority sold or bought drugs, with stimulants, cannabis, and ecstasy being the top products of choice. To this extent, these particular actors predominantly dealt in one product type but could diversify with two to three additional products. Among the countries shipped from or shipped to, the United States, United Kingdom, and Netherlands were the top three. These findings indicate that the actors most influential to the network stability of Abraxas bartered primarily in stimulants and were affiliated with the United States in some capacity. Notably, all deleted nodes across the applicable interventions were vendors.

Table 4: Descriptive Statistics for Top 100 Actors Held in Common Across Targeting Strategies

Descriptive Statistics	Degree Centrality (78)	Eccentricity (0)	Total Purchasing Price (68)	Reputation (76)	Unique Items Purchased (43)	Random (0)
<i>Actor Designation</i>						
Vendor	100% (78)	-	100% (68)	100% (76)	100% (43)	-
Buyer	0% (0)	-	0% (0)	0% (0)	0% (0)	-
<i>Number of Unique Item Categories</i>						
1	85.9% (67)	-	83.8% (57)	85.5% (65)	74.4% (32)	-
2	10.3% (8)	-	11.8% (8)	10.5% (8)	18.6% (8)	-
3	0% (0)	-	0% (0)	0% (0)	0% (0)	-
4	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
5	2.6% (2)	-	2.9% (2)	2.6% (2)	4.7% (2)	-
<i>Listing Categories</i>						
Custom Listing	7.7% (6)	-	8.8% (6)	7.9% (6)	14.0% (6)	-

Digital Goods	7.7% (6)	-	4.4% (3)	9.2% (7)	16.3% (7)	-
Drug Paraphernalia	2.6% (2)	-	4.4% (3)	2.6% (2)	4.7% (2)	-
Drugs	96.2% (75)	-	100% (68)	94.7% (72)	90.7% (39)	-
Other	3.8% (3)	-	4.4% (3)	3.9% (3)	7.0% (3)	-
Services	6.4% (5)	-	5.9% (4)	6.6% (5)	11.6% (5)	-
<i>Number of Unique Item Subcategories</i>						
1	35.9% (28)	-	35.3% (24)	34.2% (26)	18.6% (8)	-
2	23.1% (18)	-	25% (17)	23.7% (18)	18.6% (8)	-
3	23.1% (18)	-	23.5% (16)	23.7% (18)	32.6% (14)	-
4	3.9% (3)	-	2.9% (2)	4% (3)	4.7% (2)	-
5	6.4% (5)	-	5.9% (4)	6.6% (5)	11.6% (5)	-
6	2.6% (2)	-	2.9% (2)	2.6% (2)	4.7% (2)	-
7	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
8	2.6% (2)	-	1.5% (1)	2.6% (2)	4.7% (2)	-
9	0% (0)	-	0% (0)	0% (0)	0% (0)	-
10+	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
<i>Listing Subcategories</i>						
Benzos	11.5% (9)	-	10.3% (7)	11.8% (9)	14% (6)	-
Cannabis	43.6% (34)	-	47.1% (32)	44.7% (34)	48.8% (21)	-
Data	3.8% (3)	-	2.9% (2)	5.3% (4)	9.3% (4)	-
Dissociatives	10.3% (8)	-	10.3% (7)	10.5% (8)	14% (6)	-
Drugs	2.6% (2)	-	0% (0)	2.6% (2)	4.7% (2)	-
Drugs Paraphernalia	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
E-Books	5.1% (4)	-	2.9% (2)	6.6% (5)	11.6% (5)	-
Ecstasy	34.6% (27)	-	33.8% (23)	34.2% (26)	46.5% (20)	-
Electronics	2.6% (2)	-	2.9% (2)	2.6% (2)	4.7% (2)	-
Erotica	3.8% (3)	-	1.5% (1)	3.9% (3)	7% (3)	-
Fraud	2.6% (2)	-	1.5% (1)	2.6% (2)	4.7% (2)	-
Hacking	2.6% (2)	-	1.5% (1)	3.9% (3)	7% (3)	-
IDs and Passports	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
Information	5.1% (4)	-	2.9% (2)	5.3% (4)	9.3% (4)	-
Miscellaneous	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
Money	2.6% (2)	-	1.5% (1)	2.6% (2)	4.7% (2)	-
N/A	12.8% (10)	-	13.2% (9)	13.2% (10)	20.9% (9)	-
Opioids	24.4% (19)	-	25% (17)	25.0% (19)	25.6% (11)	-
Other	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
Prescription	11.5% (9)	-	11.8% (8)	11.8% (9)	14% (4)	-
Psychedelics	15.4% (12)	-	14.7% (10)	13.2% (10)	20.9% (9)	-
RCs	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
Security	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
Software	2.6% (2)	-	1.5% (1)	2.6% (2)	4.7% (2)	-

Steroids	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
Stimulants	52.6% (41)	-	58.8% (40)	51.3% (39)	51.2% (22)	-
Weapons	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
<i>Number of Unique Locations Shipped From</i>						
1	76.9% (60)	-	79.4% (54)	76.3% (58)	67.4% (29)	-
2	18% (14)	-	14.7% (10)	18.4% (14)	23.3% (10)	-
3	3.9% (3)	-	4.4% (3)	4% (3)	7% (3)	-
4	1.3% (1)	-	1.5% (1)	1.3% (1)	2.3% (1)	-
<i>Locations Shipped From</i>						
Australia	9% (7)	-	11.8% (8)	9.2%	7.1% (3)	-
Belgium	2.6% (2)	-	1.5% (1)	2.6%	-	-
Belize	1.3% (1)	-	1.5% (1)	1.3% (1)	2.4% (1)	-
Bulgaria	1.3% (1)	-	1.5% (1)	1.3% (1)	2.4% (1)	-
Canada	1.3% (1)	-	1.5% (1)	1.3% (1)	-	-
Denmark	1.3% (1)	-	1.5% (1)	1.3% (1)	-	-
Europe/EU	11.5% (9)	-	11.8% (8)	13.2% (10)	21.4% (9)	-
France	2.6% (2)	-	2.9% (2)	2.6% (2)	2.4% (1)	-
Germany	21.8% (17)	-	22.1% (15)	21.1% (16)	21.4% (9)	-
India	1.3% (1)	-	1.5% (1)	1.3% (1)	2.4% (1)	-
Italy	1.3% (1)	-	1.5% (1)	1.3% (1)	2.4% (1)	-
Netherlands	17.9% (14)	-	17.6% (12)	17.1% (13)	19% (8)	-
Norway	-	-	1.5% (1)	-	-	-
Spain	2.6% (2)	-	1.5% (1)	2.6% (2)	4.8% (2)	-
UK	19.2% (15)	-	17.6% (12)	18.4% (14)	14.3% (6)	-
United States	19.2% (15)	-	22.1% (15)	19.7% (15)	26.2% (11)	-
Unknown or N/A	14.1% (11)	-	8.8% (6)	14.1% (11)	21.4% (9)	-

Discussion

Adaptive sequential node deletion was applied to test the efficacy of six law enforcement strategies in disrupting the ease of operation of the Abraxas cryptomarket. In a real-world setting, this mimics law enforcement efforts to target, apprehend, and/or arrest individual actors within a network's value chain. As it is apparent from the results, random targeting was found to be the least effective strategy across the five outcome measures, producing minimal disruptive effect at a relatively slow pace. These results are consistent with findings from other criminal network research (Bright, Greenhill, and Levenkova, 2010; Keegan et. al, 2010; Westlake, Bouchard, and Frank, 2011) which found that "random interventions perform more poorly compared with strategies that target actors".

Curiously, random targeting was not the poorest performing strategy as eccentricity targeting proved to be the least disruptive. It is not clear why this is the case as eccentricity measures the distance of one node from every other node in a network and would presumably

prove effective as a calibrated intervention. Yet, from closer examination, it appears that the eccentricity values among the nodes were so similar such that this intervention provided little strategic value as it targeted nodes that yielded negligible disruptive impact. In other words, there were no power law distribution in the eccentricity scores such that influential nodes could be removed from the network. This naturally resulted in the low impact node removals. Nevertheless, degree centrality and reputation targeting were the most effective strategies across all five outcome measures, consistently producing near-identical results.

It is highly likely that these strategies are interrelated as the specific actors that are targeted are the same or similar. Degree centrality can be operationalized as the total number of unique buyers a vendor has done business with and vice-versa. The size of a vendor's clientele list is indicative of a more broad-based form of trust. On the other hand, reputation targeting is the preeminent marker of trust on cryptomarkets. To be clear, the more trading partners an actor has the more likely it is that they possess an equally high cumulative reputation score (Christin, 2013; Decary Hetu et. al, 2017).

While total purchasing price targeting and unique items bought/sold targeting were not quite as effective as degree centrality and reputation targeting, they did provide comparable levels of disruption to the transactional network. The disruption pattern demonstrated by these four targeting strategies was as such: the proportion of potential measurable values increases or decreases as more actors are removed. These values plateau as the network becomes completely fragmented. The disruptive effect begins to decline as the network size decreases due to vertex deletion.

Targeting Based on Human and Network Capital

This raises an interesting question about the underlying differences between these four targeting strategies. In short, are these strategies one in the same given their comparable patterns of disruption? Similarities between disruptive impact can be attributed more to the sameness of the actors targeted than to the idiosyncrasies of the targeting strategies. To elaborate, an actor that has a high degree centrality and an equally high cumulative reputation are also likely to have a high market share (revenue generated) and number of unique items bartered for. In fact, the majority (> 50%) of deleted nodes were held in common among degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting for nearly all outcome measures (see table 3). High impact nodes are likely to dominate a market across a number of metrics tied to human and social capital.

To this extent, it is evident that network and human capital metrics may not differ completely if the actors that are removed are the same. Contrary to previous studies (Bright, Greenhill, and Levenkova, 2010; Bright, Greenhill and Morselli, 2014; Tsvetovat and Carley; 2003), I posit that targeting criminal actors based on a human capital approach may not always be accurate. In short, if the network and human capital measures are interrelated or correlated to some extent, preferencing one approach over the other is functionally moot as both approaches achieve similar or perhaps near-identical disruptive impact.

However, there is one element of the human capital approach which stands out: the role of the actor. Similar to Duijn et al. (2014), I found that vendor deletion exclusively produced outsized impact on the transactional network (see table 4). Moreover, the removal of buyers was ineffectual as they were merely customers that did not supply illegal contraband. As such, their reduced engagement with actors on Abraxas precluded their prioritization by any targeting strategy aside from eccentricity and random targeting. While buyers might move the market, determining the ebb and flow of economic transactions as

they purchase the products that are advertised, it is vendors that ultimately supply these products. As such, cryptomarket interventions should be vendor-centric.

Aside from the role of the actor, it is evident that the top deleted nodes are product specialists that are based in Western nations like the United States and United Kingdom (see Table 4). In fact, the popularity of particular goods (stimulants and marijuana) and the countries from which they are shipped gives us an idea of what there is a demand for and where that demand comes from. As it relates to the general distribution of products and countries on Abraxas (see Appendix), cannabis (34.21%), stimulants (19.38%), ecstasy (13.8%), opioids (10.8%), and psychedelics (6.75%) account for the top five products sold while Germany, the United States, the UK, the Netherlands, and Australia accounted for 25.1%, 19.34%, 13.78%, 9.22%, and 8.74% of nations shipped from, respectively. As such, law enforcement interventions against dark web markets might just as well prioritize vendors that sell a specific product or ship from a specific country. As the majority of cryptomarket transactions are conducted by a small number of product and country-specific vendors, it may be beneficial to calibrate interventions based on this. While most dark web markets will sell a wide assortment of products that are shipped from a wide variety of nations, it is evident that most transactions involve a small number of product types from a short list of countries.

Metagames and Power Laws

Consistent with research in criminal networks (Druxbury and Haynie, 2019; Wood, 2017), I find that removing the most prolific vendors in sequential order fragments the network in relatively little time. Indeed, these results are not altogether different from Duxbury and Haynie (2020) who documented the existence of a scale-free online drug market and distribution network. To this extent, disassortative mixing in Abraxas, pursuant to preferential attachment, while necessary for successful transactions at scale produced clear vulnerabilities in the network structure of this market. However, this was only observed for degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting. In short, Abraxas is comprised of a small number of influential actors whose deletion would result in a fragmentation cascade. Importantly, the removal of the top nodes across these four targeting strategies yielded cross-cutting impact, producing noticeable disruption across all five outcome measures. This is particularly noteworthy as it implies that the removal of prolific actors has universal disruptive impact on the transactional network.

Furthermore, the findings indicate that when interventions are successful the disruption abides by a power law where a small number of deleted nodes produces an outsized portion of the disruption impact (see figures 2a-e). Importantly, this study establishes differences in linear, curvilinear, and power law disruption. A linear relationship means that node deletion and disruption impact are proportional, implying that the removal of a large number of nodes will result in an equally large level of disruption. A curvilinear relationship implies that a moderate number of deleted nodes account for a large amount of disruption impact. Finally, a power curve implies that small number of deleted nodes accounts for an outsized portion of disruption impact. Across the five outcome measures, degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting all demonstrated power law properties whereas eccentricity and random targeting generally demonstrated linear properties but were sometimes curvilinear.

Importantly, the disruptive impact of sequential node removal can be described as a chess-like metagame. A metagame presupposes that there are underlying rules within a game

such that understanding and abiding by them confers strategic dominance over those who understand and abide by baseline rules. In chess, the metagame involves anticipating what moves one's opponent might make and making moves which manoeuvres one's opponent into a position favourable to one's self. Moreover, one might make certain moves which set up successive moves which have a greater impact down the line.

Sequential node removal has a similar metagame where disruptive impact can be maximized if certain nodes are first removed to make way for the removal of other nodes. In short, initial nodes must be removed in order for maximal impact to be achieved once the network is reformed following the initial intervention. When examining the disruption impact, it is evident that the nodes which produced the greatest impact were often those that were not first removed (i.e. had the highest value per the parameters of a specific targeting strategy). In fact, nodes which had the greatest disruptive impact were often those outside of the first 10 nodes that were deleted. As it relates to law enforcement interventions, initial arrests or apprehensions should be used to set up future arrests or apprehensions that have a greater capacity for disrupting the criminal network.

Conclusion

Adaptive computer simulations represent a novel means of testing the structural robustness of a criminal network as well as the effectiveness of strategic interventions. However, the results of these analyses are driven by pre-determined parameters which govern the behaviour of the actors within the network. While driven by educated and evidence-based suppositions, these parameters are fundamentally speculative. As such, these results should not be mistaken for actionable intelligence gathered from a real-world experiment. They can only go so far in explicating the true dynamics which undergird the phenomenon of study. Randomized control trials represent the gold standard of research within the social sciences. In this regard, adaptive computer simulations are a secondary option. Future research into cryptomarket disruption strategies should consider experiments on live markets where interventions are attempted, and the results are measured. While the logistical difficulties of such an undertaking are understandably large, such an experiment represents the pinnacle of evidence-based research into cryptomarkets.

Conclusion

Based on the various analyses conducted in this thesis, it is evident that cryptomarkets offer unique opportunities for researchers seeking to examine the macro-level structure of an illicit transactional network as well as the functional mechanisms which undergird consumer activity therein. To reiterate, this dissertation had two overarching objectives. First, it sought to push the theoretical boundaries of cryptomarket research in order to better understand the functional mechanisms of cryptomarkets. To specify, I sought to identify network dynamics between Abraxas participants pursuant to the formation of trust, the predictors of consumer satisfaction based on consumer feedback, and the efficacy of strategic interventions against cryptomarkets. The second aim of this dissertation was to leverage the various findings therein to inform targeted interventions by law enforcement against cryptomarkets. To this extent, the task of understanding which strategies work and which ones do not hinges on a more fluid understanding of a phenomenon.

While Abraxas is by no stretch of the imagination representative of all cryptomarkets given its size and relative influence, it is nevertheless a platform which possesses all of the characteristics and qualities of a standard dark web market. As a result, while the results and conclusions drawn from these chapters are not perfectly generalizable to all cryptomarkets, they should serve to instruct law enforcement activity on the dark web. To this extent, analyses of the Abraxas trade network have provided key insight into three important areas of inquiry.

As chapter one of this thesis functioned as a large-scale literature review, no particular practical insight can be drawn from it. Nevertheless, this chapter served to illustrate the state of the scholarly literature on cryptomarkets and the areas where more research is needed. Importantly, these specific areas were covered in the successive chapters of this thesis. Following Duxbury and Haynie (2017), chapter two examined the network structure of Abraxas in order to identify the market-level metrics that predicted for vendor selection. These findings provide more insight into how trust among buyers and vendors determines the structure of a cryptomarket. In particular, this chapter sought to test the generalizability of Duxbury and Haynie's (2017) initial study to determine if preferential attachment pursuant to trust dynamics played a role in the topology of a cryptomarket. Furthermore, this study offered insight into the predictors and the development trajectory of vendor trustworthiness.

Chapter three identified and compared the determinants of customer satisfaction and dissatisfaction among buyers. It, moreover, was the first study to apply text mining methods to determine whether the sentiment structure of qualitative reviews differed between five-star and non-five-star ratings as well as transactions that were finalized early and those that were not. In short, this chapter sought to illuminate lexical features associated with consumer satisfaction and dissatisfaction on a cryptomarket. The theoretical thrust of this chapter centred on information asymmetry and its role in allowing buyers to make an informed decision on which vendors to do business with. Information asymmetry is a particularly important concept as it helps unravel the inner workings of trust dynamics as well as the network structure of cryptomarkets. For law enforcement operations against cryptomarkets to have a long-lasting impact, there should be a sombre consideration of how each operation can facilitate more information asymmetry in the cryptomarket environment.

Adding to Duxbury and Haynie's (2020) study, chapter four examined how sequential node deletion may affect a cryptomarket's ease of operation. To this extent, computer simulations which incorporated network adaption and preferential selection were leveraged to better understand which strategic intervention(s) were most effective at disrupting the structural integrity of Abraxas' network structure. This particular study is important for cryptomarket disruption strategies as it demonstrates that the behaviour of an illicit trade network can be modelled (Duxbury and Haynie, 2019) and subsequently vivisected through an evidence-based calculus. Moreover, it provides insight into how law enforcement might approach the curtailment of a cryptomarket. As cryptomarket takedowns and the opportunistic arrest of vendors are not particularly effective in the long-term disruption of these entities, a carefully calibrated intervention which considers network dynamics such as preferential selection is warranted.

Summary of Findings

Chapter 2

Like Duxbury and Haynie (2017) and Norbutas (2018), social network analysis of the Abraxas cryptomarket revealed a large and diffuse network where the majority of buyers purchased from a small cohort of vendors. The transactional network is quite diffuse with a network density of 0.0007. Furthermore, the full network consisted of 29 components, with the largest component containing 97.6% (2726) of all nodes within the network. The remaining connected components consisted of 19 dyads, 7 triads, and single assortments of components of various sizes. In short, Abraxas buyers tended to purchase from a small number of vendors over time, which leads to the formation of a large group of sparsely connected users with very few isolated buyer-seller cliques. This theme of preferential selection of vendors on the part of buyers is repeated in other findings within this study. More generally, this study provides further evidence of the role of preferential attachment in the network structure of dark web markets.

The average out-degree centrality was 2.16 whereas the average in-degree centrality was 20.2. This indicates that buyers did business with 2.15 vendors, on average, while vendors did business with an average of 20.2 buyers. This pattern can be gleaned from the community detection analysis. This revealed 158 unique communities formed around prolific vendors. The largest of these communities possessed 390 members, whereas the smallest 111 communities have fewer than 10. Not surprisingly, 35 and 20 communities were dyads and triads, respectively. What is particularly telling is that the leading 20 communities accounted for 63% (1763) of all actors and 71.9% (3909) of all transactions. Moreover, the average community had 1.7 vendors and 15.98 buyers. In other words, each vendor and their respective buyers constitute individual communities. The Abraxas transactional network can then be viewed as set of transactional islands as opposed to a large, densely connected conglomeration of vendors and buyers. It is also important to note that these transactional communities within the network were country and product-specific, meaning that a specific product type was shipped to a single country.

Regression analyses for vendor success, popularity, and affluence demonstrated that the cumulative reputation score of vendors was the predominant predictor for trust across all three proxy variables. Additionally, cumulative risk was the second statistically significant predictor across all three models. This indicates that a vendor's willingness to incur the risks associated with overseas shipping yields greater economic opportunities and with it a

reputation for trustworthy conduct. This corroborates Decary-Hetu's (2016) contention that cryptomarkets facilitate localized trading where buyers from a specific nation purchased from vendors who also ship to that specific nation. This is done to mitigate the risks associated with overseas shipping. However, each model differed on what particular estimates explained the variance in vendor success, popularity, and affluence. While cumulative purchase price, item categories, and item subcategories were positive predictors of vendor success, unique item listings and subcategories were predictors for vendor popularity and the average purchase price and number of words in the item description predicted vendor affluence. As such, it is evident that the price and the type of product sold influences how successful a vendor is relative to others operating on the same market.

Finally, the results garnered from the social network analysis are corroborated by the trajectory analyses. Indeed, the trajectory models demonstrated that a small number of vendors become highly successful, popular, and affluent in a relative short period of time. Moreover, vendors that possess a specific ranking within the market will likely remain as such throughout the market's operation as low-achieving will remain low-achieving while high-achieving vendors will become increasingly successful, popular, and affluent on the market. This is reflective of the law of accumulated advantage whereby those with many resources continue to gain more resources.

This lends credence to the idea that buyers' preference for a vendor has as much to do with the specific product and price requirements of buyers themselves than it does with the trustworthiness of a vendor. In short, buyers will select vendors with whom they trust but who are also providing the products they desire at the prices they can afford. For example, a buyer seeking to purchase marijuana will have no dealings with a vendor trafficking in counterfeit coins, regardless of their reputation. This particular bit of information should serve to inform future law enforcement interventions. Indeed, a law enforcement organization looking to destabilize a cryptomarket should first consider which products are popular on the market and which vendors are providing them. If the goal is general disruption, it makes little sense to target a vendor who traffics in a product that is not popular among buyers. For law enforcement organizations, this is the equivalent of knowing one's audience.

Chapter 3

At the outset, it was evident that Abraxas feedback was governed by the Pollyanna principle where 96% (4683) of all purchases received a rating of five. As such, it would seem that Abraxas buyers were overwhelmingly satisfied with their transactions. Nevertheless, text mining revealed acute similarities and differences between the words used to describe five-rated and non-five-rated transactions. Based on frequency analyses, "finalize" and "early" were the two most popular words among non-five-rated feedback and the second and third most popular words among five-rated feedback. However, the five-rated corpus possessed words with a positive connotation ("good", "great", "fast", "stealth", "thanks", "quality", "trust", "best", "top", and "nice") whereas words in the non-five-rating corpus were more negative or value-neutral ("update", "scam", "product", "nothing", "never", "still", and "waiting"). This suggests lexical differences between five and non-five feedback.

Not surprisingly, word associations for five and non-five-rated transactions revealed that "finalize" and "early" were highly associated with one another, but also co-occurred with "trust", "hope", and "confidence". Nevertheless, five-rated words associations are characterized by trust, satisfaction, and praise for the vendor, product, and process whereas non-five-rated word associations were predominantly value-neutral but were sometimes

negative when the transaction was described. This is also evident from the sentiment analysis. The five-rated corpus possessed a positive average polarity of 0.6 while the non-five-rated corpus was slightly negative with an average polarity of -0.01. This suggests that buyers, while dissatisfied, will not harshly criticize vendors or vent their frustrations when providing feedback to the rest of the market.

Finally, the logistic lasso regression demonstrated that “fast”, “great”, “thank”, “good”, and “vendor” predicted a five rating while “product”, “finalize”, and “early” were more likely to be associated with a non-five rating. Furthermore, the odds of a five rating are 130% and 115% greater when the words “great” and “thank” are found in the feedback than when they are not. In contrast, the presence of “product”, “finalize”, and “early” in buyer feedback increased the odds of a non-five rating by 2.1%, 3.1%, and 27.1%, respectively.

As it evident from these findings, buyers are generally pleased with their transactions on Abraxas as long as the product arrives on time and is as advertised. In general, vendors have a relatively low bar to achieve when it comes to satisfying their customers. This is perhaps due to the fact that goods and services are transparently advertised on a cryptomarket. To elaborate, because buyers are made aware of the quality of the product from the advertisement, their satisfaction is earned if the product simply meets their expectations based on the advertisement. Only in rare circumstances do buyers award a vendor a rating below a 5. This generally occurs when expectations are not met. Intuitively, this suggests that vendors can become popular and successful if they merely keep their word to buyers. As trust is scarce in any criminal setting, the factors leading to its attainment are perhaps less strenuous if all vendors need do is to keep their word. This is perhaps why popular vendors can increase the prices of their products despite the quality of their products remaining the same. Buyers are perhaps accepting of price hikes as they understand that the rarity of a trustworthy vendor precludes paying a higher price for their services. This is similar in licit economies where buyers will pay exorbitant prices for brand name products as their quality is assured.

Chapter 4

Based on the results of the sequential node deletion, random targeting was found to be ineffective across the five outcome measures, producing minimal and a slow disruptive effect. This result was not altogether unexpected. However, eccentricity targeting proved to be the least disruptive. Degree centrality and reputation targeting were the most effective strategies across all five outcome measures, consistently producing near-identical results. While total purchasing price targeting and unique items bought/sold targeting were not quite as effective as degree centrality and reputation targeting, they did provide comparable levels of disruption to the transactional network. The disruption pattern demonstrated by these four targeting strategies was as such: the proportion of potential measurable values increases or decreases as more actors are removed. These values plateau as the network becomes completely fragmented.

Furthermore, it is highly likely that these strategies are interrelated as the specific actors that are targeted are the same or similar. This suggests that while the stated objectives of these targeting strategies are different, their functional performance is the same or similar. To this extent, the network and human capital approaches pursued in this study are moot as the results were the same when sequential node deletion was undertaken. Importantly, this suggests that the most dominant vendors on Abraxas are universally dominant across a number of indicators which measure their success on the market. Finally, from analyses of

the disruptive impact of each targeting strategy, degree centrality targeting, reputation targeting, total purchase price targeting, and unique items bought/sold targeting are each based on a power law where a small percentage of deleted nodes is responsible for an outsized proportion of the disruptive impact across all five outcome measurements (e.g. 1-5% of deleted nodes were responsible for 45-90% of disruptive impact). Random and eccentricity targeting appear to be linear across four of the five outcome measurements, suggesting minimal disruptive impact.

It is clear from these findings that the power law dynamics of Abraxas makes the market susceptible to targeted attacks. To this extent, law enforcement agencies need not target the entirety of the transactional network but should instead focus on the most influential vendors as they are involved in the majority of market activity. This fits an evidence-based policing calculus where the goal is to do more with less. In this case, the ease of operation of a cryptomarket may be effectively disrupted should law enforcement officials focus on the power few vendors.

Implications for Dark Web Interventions

Based in the available evidence addressed in chapter one, it appears that law enforcement interventions against cryptomarkets have been ineffective and perhaps counterproductive. In the aftermath of market closure, sales volumes generally returned to comparable pre-closure levels while new markets emerged to take the place of those shut down. In fact, the FBI's effort to shut down the original Silk Road utterly fragmented the composition of the cryptomarket landscape. Indeed, a once consolidated market dominated by the Silk Road devolved into a hypercompetitive affair between various smaller cryptomarkets vying for volatile market shares. These new cryptomarkets included Agora, Silk Road 2.0, Black Market Reloaded, Sheep Marketplace and Pandora. Of course, Silk Road has itself undergone several resurrections, returning as Silk Road 2.0, Silk Road 3 Reloaded, and the latest iteration Silk Road Reloaded.

It is important to note that this fragmentation is partially due to the decentralized exchange networks of cryptomarkets. "In the event that a cryptomarket is shut down, the user community is able to persist; users either migrate to other sites or, as in the case of Silk Road 1.0, they construct and quickly repopulate a replacement website" (Martin, 2014a, 23). This mobility and durability equate to a difficult-to-exterminate illicit entity. Though it is perhaps reprehensible to allow the unabated operation of organized crime, it is arguably far worse for law enforcement to destroy a criminal monopoly. We saw this again in 2014 and 2017 when law enforcement shut down some of the largest cryptomarkets in operation at the time.

While it is imminently clear that largescale market closures are not the way forward, a broad-based focus on the trust dynamics within cryptomarket transactional networks may be. Based on the findings of this thesis, a carefully calibrated network-based approach which targets trusted, high-earning vendors may yield the most disruptive impact. Moreover, a strategy which randomly targets actors on a cryptomarket is not advisable from these findings. While controversial, a vendor-centric targeting strategy which exploits trust dynamics in the transactional network but leaves the market's transactional structure intact would offer maximal disruptive impact without displacing actors to rival and/or new cryptomarkets. This particular strategy resembles wild animal population control where animal populations are kept at a manageable level so as to prevent harm to humans. In a

similar fashion, cryptomarkets are permitted to function, though at a heavily reduced capacity.

Based on my findings, trust on Abraxas is predicated on a pareto distribution where a small number of trusted vendors reap the rewards of their reputation. To this extent, reputations serve as a tool for identifying the quality of merchandise and, to an extent, counteract uncertainty within a highly volatile environment. More importantly, however, trust played a key role in determining the transactional network of Abraxas. The global network structure of Abraxas is a product of initial and repeated transactions between buyers and vendors. Each vendor and their respective buyers constitute an individual community within Abraxas. These communities were also locational and product-specific, suggesting the importance of geographic distance and niche markets in moulding the network structure. This can be exploited by law enforcement.

While trust between buyers and vendors fundamentally determines the structure of a transactional network it can also be exploited to undo this structure. Trust, in other words, operates as a double-edged sword as it allows buyers to identify top vendors and law enforcement to identify high priority targets. This reveals a game theoretic problem within cryptomarkets. When buyers attempt to mitigate risk by trading with the most trustworthy vendors this creates easily exploitable vulnerabilities in the network structure of the market.

From the results in chapters two and four, the removal of vendors with the highest cumulative reputation scores (i.e. the most trusted vendors) yielded the largest disruptive impact to Abraxas. Importantly, actors with the highest cumulative reputation are also the actors with most trade partners, products sold, and revenue generated. Based on my findings, a targeting strategy which sequentially removed these prolific actors would likely result in a fragmentation cascade. Bereft of their primary vendor, buyers would presumably take their chances with a vendor they have little experience with or leave the market entirely. The scale and profitability of a cryptomarket might be curtailed by such a strategic intervention.

As Duxbury and Haynie (2020) note, “When networks are attacked, actors grow more cautious about forging ties, connecting less frequently and only to trustworthy alters.” In short, the entire premise of such a targeted intervention against a cryptomarket would be to rattle the trust and confidence of those operating on the market. When the most trustworthy operators are taken off the board the overall level of trust within the market dissipates. In the abstract, the objective is not to target vendors but to target trust.

In general, the scale free properties of Abraxas and the cryptomarket examined by Duxbury and Haynie (2017) suggests that these network topologies are premised on preferential attachment. As such, law enforcement organizations need not launch a large-scale attack on the market, targeting the entirety of the vendor cohort. Instead, maximal gains can be achieved by focusing on the power few vendors who account for the majority of sales made, buyers transacted with, and revenue generated. The premise of evidence-based policing is doing more with less and this seems a sensible option for law enforcement organizations working with scarce resources.

It is, however, important to note that there is a high-level metagame embedded within this strategy. As discussed in chapter four, a metagame assumes that there are underlying

rules within a game such that understanding and abiding by them confers strategic dominance over those who understand and abide by baseline rules. In chess, the metagame involves anticipating the opponent's probably move set and making counter moves which positions the opponent into a favourable position. Based on my findings, disruptive impact on cryptomarkets can be maximized if certain actors are first removed in order to set up higher impact removals. When examining the disruption impact in chapter four, it became evident that the nodes which produced the greatest impact were often those that were not first removed (i.e. had the highest value per the parameters of a specific targeting strategy). In fact, nodes which had the greatest disruptive impact were often those outside of the first 10 nodes that were deleted. For law enforcement agencies, initial arrests or apprehensions of cryptomarket vendors should be used to set up future arrests or apprehensions that have a greater capacity for disrupting the criminal network.

In general, metagame dynamics must be consciously considered by law enforcement organizations dealing with online and offline criminal organizations. The possibility of iatrogenic and backfiring effects must be carefully considered before an operation is launched. On the other hand, law enforcement must also consider the strategic value of an intervention at the macro-level. In other words, how might the disruption of one criminal entity affect the entire criminal ecosystem within which that entity resides? This was not considered by law enforcement who shut down cryptomarkets in 2011, 2014, and 2017. Based on the findings of this thesis, the optimal metagame strategy would involve the curtailment of existing cryptomarkets through the targeted removal of vendors as opposed to completely shutting these markets down. In summation, the targeting of key actors within a cryptomarkets serves to control the spread of market as opposed to completely eliminating it. Outright market seizure or elimination compounds the problem, creating larger and more sophisticated markets for which more resources will be required to police.

It is, moreover, an open question as to whether prior interventions against cryptomarkets were simply ineffective or if the cryptomarket environment is antifragile, growing more robust with each major shock it suffers. If the latter supposition is correct, it would make little sense for law enforcement to pursue future operations which seek to dismantle these markets in their entirety (i.e. market seizure). Rather, law enforcement resources would be better spent in targeted (or pinprick) interventions which curtail the growth these markets through the removal of prolific actors that drive market activity. Indeed, the embeddedness of cryptomarkets may mean that these illicit entities are incapable of being eradicated in their entirety. Nevertheless, such a strategy would also rely on entropy within the cryptomarket environment. As markets are generally operational for several months (Christin, 2013) and are subject to closure due to the duplicity of the actors therein, law enforcement may seek to play to this dynamic when targeting prolific actors on large markets. To this extent, this might involve leaving small and uninfluential markets to their own devices while targeting actors generating the most active vendors on on the largest actors. In allegorical terms, this strategy equates to catching the biggest fish in the largest pond while leaving smaller fish to die as smaller ponds dry up.

This notion of "leaving cryptomarkets to their own devices" is a particularly controversial decision as it implies that law enforcement organizations are simply allowing criminal groups to operate unimpeded. This notion is, however, incorrect. As resources are

limited within a policing context, not all crime and criminals can or should be policed equally. Indeed, criminals and crime events are not equal in the damage they cause or the resources that are required in order to adequately police them. Such is the aim of strategies like hotspots policing or targeted foot patrols which target offline crime. These particular strategies prioritize the areas most afflicted by crime, allocating resources to locales which need them most. A similar logic can be applied to the cryptomarket intervention strategy I have proposed. The overarching goal is to optimize the resources expended. In other words, such a strategy intends to get the most bang for one's buck, targeting areas of the dark web environment where the largest possible impact can be made without compounding the problem further.

Nevertheless, something must be said about the potential for displacement in the face of targeted interventions. Based on the criminological literature (Johnson, Guerette, and Bowers, 2014), there is scant evidence of widespread geographic crime displacement following targeted patrols. To this extent, 30 years of research on this topic suggests that crime is displaced in only a small number of cases. There is, moreover, a diffusion of crime reduction benefits where surrounding areas experienced a pronounced crime drop that was comparable to the targeted areas (Guerette and Bowers, 2009). This pattern has been acknowledged and further proven by several meta-analyses and systematic reviews (Bowers et al., 2010; Braga et al., 2012; Telep et al., 2014). Still, there is one glaring limitation with regard to research on crime displacement: the extensive focus on geography. Indeed, comparatively little is understood about other forms of displacement and diffusion. This includes temporal, target, tactical (use of method or tools), and crime type displacement (Hesseling, 1994).

While our understanding of crime displacement in physical settings is generally well-developed from a geographical perspective, this is not the case for displacement in a cyberspace. Indeed, we know very little about the vicissitudes of crime displacement in this environment, and far less about its occurrence on cryptomarkets. Unlike terrestrial environments where criminal opportunities are predicated on unique environmental characteristics, cryptomarkets are not themselves diverse entities. Rather, many, if not all, cryptomarkets possess the same infrastructure, financial risk reduction competencies, and operational practices. As such, if one market were to shut down, buyers and vendors would simply migrate to another market. In this case, crime displacement of some sort is a clear and ever-present reality in the cryptomarket environment. However, this dynamic might differ somewhat when we consider the removal of specific vendors and the resultant displacement of buyers to other vendors on the same or different markets.

To this extent, crime displacement on cryptomarkets might occur in four forms: market-based displacement, vendor-specific displacement, product-based displacement, and platform-based displacement. While market-based displacement (movement from one cryptomarket to another) is easily understood, vendor-specific displacement refers to the movement of buyers from one vendor to another following the removal or absence of a vendor on the same cryptomarket. As discussed, much of this is predicated on the development of trust where reputable vendors can serve as viable replacements to vendors that are no longer present on a cryptomarket. Similar to crime type displacement in terrestrial settings, product-specific displacement refers to the changes in the product purchasing habits

of buyers where the removal of a specific product type engenders movement toward the purchase of a different product type. For example, the widescale removal of fentanyl from cryptomarkets may encourage buyers to purchase a lower grade synthetic opioid to compensate. Finally, platform-based displacement refers to the movement of vendors and buyers away from cryptomarkets into terrestrial markets.

Given these varying forms of displacement, one cannot rule out the possibility that any proposed intervention, whether premised on computer simulations or otherwise, cannot and will not yield some level of displacement. As a conceptual matter, displacement is always a possibility in an environment where criminal opportunities are contingent on environmental factors. This is more so the case when these criminal opportunities are present on an online platform which is not subject to the same constraints levied upon offline platforms. Indeed, the strategic and incremental removal of high-value vendors from a cryptomarket might encourage to buyers to move to another vendor or market. Moreover, buyers might switch to another product or outright leave the dark web for an offline market. Moreover, given the anonymity of dark web marketplaces, it is difficult to determine the extent to which displacement may occur as actors are liable to use different profiles when engaging on different markets.

Bastions of Responsible Use

Given the aforementioned issues and questionable benefits of cryptomarket interventions, it is an open question as to whether law enforcement should target these markets at all. There is, moreover, an extensive and long-running debate regarding the harm reduction capabilities of drug enforcement. Given the negative externalities created by police crackdowns on drug markets, the goal of law enforcement may not involve the eradication of drug markets, but the reduction of the potential harm caused by the transaction and consumption of illicit substances. Cryptomarkets fit neatly within this discussion as they are, in many ways, a viable and more preferable alternative to terrestrial markets and street dealing.

Cryptomarkets serve to mitigate the negative externalities endemic to terrestrial drugs markets, namely physical violence. Indeed, violence is difficult to actuate on cryptomarkets given the immateriality of cybercrime. To this extent, the anonymity and geographical dispersion afforded to cryptomarkets means that participants cannot simply harm other actors. The improbability of violence on cryptomarkets lies in the platform's dematerialization of voluntary economic transactions. This has been documented in several studies (Aldridge and Decary-Hetu, 2014; Morselli et al., 2017; Van Hout and Bingham, 2013a). One study found that cryptomarket vendors had a smaller likelihood of experiencing violence relative to "street" dealers as most of their clientele were middle-class, university students that were averse to violence (Mohamed and Fritsvold, 2010). Furthermore, Barratt et al. (2016), surveying 3794 respondents from 57 countries on drug use, found that 1.3% and 1% of cryptomarkets users experienced "threats to personal safety" and "physical violence", respectively. In contrast, 14% and 6% of those who purchased from friends, 24% and 10% of those who purchased from known dealers, and 35% and 15% of those who purchased from strangers experienced "threats to personal safety" and "physical violence", respectively. In general, buyers reported safer and more convenient transactions given the complete

circumvention of face-to-face meetings with potentially dangerous dealers (Barratt, Lenton, and Allen, 2016; Van Hout & Bingham, 2013a, 2013b).

Martin (2018) contends that “cryptomarkets are displacing potentially violent drug market norms in favour of more cordial relationships between market participants.” Martin (2018), moreover, refers to this as the “gentrification hypothesis” whereby the safety and anonymity of illicit online transactions precludes the use of and necessity for violence. Cryptomarket vendors compete on the basis of reputation, relying on the quality of their products and marketing campaign. As Martin indicates, the cryptomarket vendors are encouraged to create a “socially constructive public image that is both free from violence and more attuned to the perceived priorities of their customer base” (2014a, 40). Creating rapport and behaving in a trustworthy manner go farther on cryptomarkets than would violence were it an option available to actors.

The importance of vendor reputations is intimately tied to the quality of goods and services offered on cryptomarkets. To this extent, consumer feedback mechanisms serve to reward the accountability of vendors. As a result, the quality of products on cryptomarkets is likely to be higher compared to offline markets. As Horton-Eddison et al. (2021, 6) contend, “this is important because some drug harms arise from uncertain content and strength, thereby creating the risk of unwanted effects or overdose.” Furthermore, vendors will often provide warning labels which inform buyers of the potential dangers of specific products. This allows buyers to make safer purchases which they could not otherwise do in a terrestrial market where street dealers are less than forthcoming about their wares. Furthermore, cryptomarkets may opt to remove or ban products that are harmful to users. Such is the case for fentanyl, assassinations, child pornography, and weapons of mass destruction.

“In addition, cryptomarket discussion forums have provided a rich source of drug safety information (e.g., quality, purity, adulterants, dosing), enabling buyers and vendors alike to share information about product and batch content, and about buying and selling more safely” (Horton-Eddison et al., 2021, 7). This information is often absent in clearnet forums much less offline markets. While the information provided on cryptomarket forums are not guaranteed to be accurate, the adoption and spread of best practices on these platforms are often hosted by qualified drug harm reduction professionals. Such is the case of Dr. Fernando Caudevilla who provided expert harm reduction advice to buyers and vendors operating on the Silk Road and other dark web markets.

While outright support for cryptomarkets by state actors is unfeasible, there is an argument to be made about the merits of toleration. To this end, state actors may choose to allow these platforms to operate as they may offset the violence and customer harm endemic to terrestrial markets. This may involve the conversion of illegal markets into licit markets. Such is the case of Portugal and some states in the United States that have legalized specific drug markets.

Theoretical Contribution to Criminology

As discussed in the first chapter of this dissertation, cybercrime, much like terrestrial crime, is the product of the intersection of three requisite factors: a suitable target, motivated offender, and lack of a capable guardian. This is referred to as routine activities theory

(Cohen and Felson, 1979). Per Cohen and Felson's (1979) theory, the infiltration of a database or distribution of malware, for example, must possess these qualities for it to have taken place (Grabosky and Smith, 2001). While some researchers (Yar, 2005) have questioned the straightforward application of routine activities theory to cybercrime, it is my contention that routine activities theory explains much about the operation of cryptomarkets and the dynamics of buyer-vendor relations therein.

When choosing to engage in voluntary economic transactions, a cryptomarket buyer must have some level of trust in a vendor's ability to make good on their promises. Indeed, trust is the key element which allows partnerships in the criminal world to form and proceed forward. While a crime must take place at a specific time, at a specific location, using specific tools, against a specific target, trust is the constituent element which must be present if the crime has more than one offending party. While logically sound in its description of criminal activity, routine activities misses one major element of criminal activity: interpersonal trust. Indeed, based on the various results of this thesis, trust pursuant to interpersonal relations among prospective criminals is a fundamentally important element which must be established for a crime to take place when more than one offender is involved. Rarely is crime an individual activity. Rather, is it an activity that is, more often than not, born out of the coordination and collaboration of multiple actors. This is certainly the case when we examine cryptomarkets.

Nevertheless, trust is a difficult-to-establish element in the criminal world. Wright and Decker (1994) and Hamill (2011) observed that betraying one's friends, family, and associates is normal in the criminal underworld. Indeed, the situational constraints with which a criminal must contend (death, arrest, betrayal, etc.) certainly encourages thoughts of renegeing on contractual obligations and turning tail when circumstances dictate. To make matters worse, these contractual obligations are not upheld by a principal authority as they would be in licit markets. Nevertheless, trust is the tool which allows criminals to cooperate, ultimately permitting the heist, assassination, or arson to move forward. It is, moreover, important to consider how one determines whether or not their fellow criminal can be trusted. Williamson (1993) maintains that this requires a trustee to demonstrate to the truster a temporary suspension of selfish desires for the sake of cooperation (458). The trust deficit within the criminal world is particularly problematic for trusters.

Crucially, trust has a curious effect on the other strands of routine activities theory. Moreover, it can also be affected by the presence or absence of these requisite characteristics. Indeed, the motivation of offenders might differ based on the trust offenders has in one another. Low trust among criminals might reduce their motivation to move forward with the crime while high levels of trust might engender greater motivation. Of course, it may also be the case that offender motivation also affects trust. To this extent, sufficiently high levels of motivation among offenders may increase the likelihood of trust between them. As such, trust and the motivation to commit crime are circularly linked with each element affecting the other.

Nevertheless, trust might also affect how offenders view the feasibility of committing a prospective crime. Indeed, the trust one puts in their fellow criminal might affect how an offender views a crime opportunity as greater trust is likely to increase the willingness to take

advantage of a crime opportunity on the part of criminals. Perhaps it is the case that trust must reach a sufficient level such that the criminal opportunity is seen as worthwhile. Of course, it may also be the case that the difficulty of completing the criminal opportunity might have an affect on the trust criminals put in one another. For example, a difficult heist might strain relations between co-offenders given the intricacy of the crime in question. Finally, the presence or absence of a capable guardian assuredly affects the trust offenders put in one another. Trust is likely to be higher in the absence of a capable guardian and lower when a capable guardian is absent.

While the criminological literature (Gambetta, 2000; von Lampe and Johansen, 2004; Gambetta, 2009; Campana and Varese, 2013) has emphasized the trust deficit within the criminal world, these observations reflect criminal activities which take place in terrestrial markets. However, based on the findings of this dissertation, these trust dynamics are not altogether different in cyberspace. As a theoretical matter, trust follows a power law distribution on Abraxas and, in all likelihood, on other criminal markets. To this extent, there is a suffusion of trust in a small number of cryptomarket vendors. This raises another question: is trust a finite commodity in criminal environments?

While this question cannot be precisely answered without ethnographic data, the presence of a power law on Abraxas serves as circumstantial evidence. Indeed, a small number of vendors become highly successful, popular, and affluent in a relative short period of time. In other words, trust is disproportionately concentrated in a small number of vendors who reap the rewards. In this case, it seems likely that trust on Abraxas is predicated on a “winner-take-all” schema where select vendors who are able to attain the trust of buyers come to dominate the market throughout its operation. This is evident in the network structure of Abraxas as well the effects of the sequential removal of nodes.

When choosing to engage in crime, criminals working with a partner must evaluate the criminal opportunity, the presence or absence of a capable guardian, and their own motivation. However, each of these elements is contingent on the trust they bestow upon their partner. From this perspective, if offenders do evaluate the risks and rewards associated with the commission of a crime, they must also consider how their partners might also perceive these risks and rewards. The calculus is further complicated by the addition of more criminal actors, with each additional actor creating new considerations for all involved.

Future Research

Cryptomarkets represents a fascinating area of study for researchers interested in the intersection of cybercrime and network science. Indeed, these platforms present a novel opportunity for researchers to test the accuracy of key theoretical precepts that are present in terrestrial markets.

The studies featured in this thesis had the explicit aim of either examining sparsely researched or entirely unresearched topics in the cryptomarket scholarship. Given the dearth of research on the network structure and resilience of cryptomarkets and the determinants of consumer satisfaction among cryptomarket buyers, these are all areas where more research is required. As such, it is suggested that future research continue to examine these particular topics, testing their generalizability on other markets. There are, moreover, a number of different methodological approaches that might be pursued in these future studies. It is likely that the results produced might be slightly or entirely different had different techniques been

used. As such, researchers should endeavour to push the methodological boundaries of cryptomarket research, developing standard procedures by which data can be more efficiently analysed.

Aside from these topics, there are number of other areas where more research is needed. According to Barratt and Aldridge (2016), “we do not yet have good evidence to indicate what proportion of the population may be sourcing drugs from cryptomarkets, and whether their numbers may be increasing” (9). Given the increasing technological sophistication of younger generations, it is also an open question as to whether cryptomarkets are primarily frequented by those defined as millennials and gen z. Moreover, it is unclear why these individuals choose to purchase drugs and other illegal goods and service on the dark web as opposed to or in tandem with terrestrial markets.

In this regard, there are several pressing questions which must be asked and potentially answered by cryptomarket scholars. What is role of cryptomarkets in facilitating new trends in drug use? To this extent, what role, if any, have cryptomarkets played in the proliferation of fentanyl and other synthetic opioids? What is the demographic profile of those who set up and operate cryptomarkets? While administrators such as Ross Ulbricht have been arrested, it is unclear who exactly establishes cryptomarkets and, more importantly, what their motivations are. At a macro-level, how does migration from terrestrial markets to cryptomarkets affect the incidence of violence as well as the wellbeing of cryptomarket participants? While it is clear from the literature that cryptomarkets reduce violence, it is unclear how much violence is potentially reduced as a result. This bears political implications as the widespread use of cryptomarkets may engender calls for further drug legalization. How have cryptomarkets innovated in response to law enforcement interventions? How fast were these adaptations made and how effective have they been? This particular set of question deals with the innovative nature of cryptomarkets, an area which may aid law enforcement in understanding the potential outcomes of future interventions.

These are some of the more pressing questions which should be answered by researchers examining cryptomarkets. Nevertheless, this is not an exhaustive list as more questions abound. Nevertheless, cryptomarkets represent a potentially worthwhile area for criminologists to research. Given the increasing technologization of crime, it is one criminological phenomenon which bares serious implications for the future of illicit trade.

Appendix (Chapter 4)

<i>Descriptive Statistics on the Abraxas Cryptomarket</i>		
Descriptive Statistics	Mean (SD) or Total	Range
<i>Vendor Reputation</i>		
Cumulative Reputation	98.76 (191.46)	0-1628
Average Reputation	4.85 (0.54)	0-5
Cumulative Positive Reputation	97.43 (189.7)	0-1625
Cumulative Negative Reputation	1.327 (4.67)	0-59
<i>Ratings</i>		
0	1.4% (74)	-
1	0.4% (23)	-
2	0.2% (10)	-
3	0.5% (26)	-
4	1.1% (59)	-
5	96.5% (5242)	-
<i>Listing Categories</i>		
Drugs	92.9% (5050)	-
Digital Goods	5.9% (321)	-
Services	0.4% (21)	-
Drug Paraphernalia	0.3% (17)	-
Other	0.3% (14)	-
Custom Listing	0.2% (11)	-
<i>Listing Subcategories</i>		
Cannabis	34.21% (1859)	-
Stimulants	19.38% (1053)	-
Ecstasy	13.8% (750)	-
Opioids	10.8% (587)	-
Psychedelics	6.75% (367)	-
Benzos	3.7% (201)	-
N/A	2.72% (148)	-
Prescription	2.19% (119)	-
Dissociatives	1.25% (68)	-
Information	1.03% (56)	-
E-Books	0.98% (53)	-
Erotica	0.9% (49)	-
Fraud	0.59% (32)	-
Steroids	0.35% (19)	-
RCs	0.22% (12)	-
Data	0.2% (11)	-
Drugs (Cyber)	0.17% (9)	-
Hacking	0.15% (8)	-
Money	0.11% (6)	-
Weapons	0.11% (6)	-

Electronics	0.09% (5)	-
IDs and Passports	0.07% (4)	-
Other	0.06% (3)	-
Software	0.06% (3)	-
Miscellaneous	0.04% (2)	-
Security	0.04% (2)	-
Drugs Paraphernalia	0.02% (1)	-
Services	0.02% (1)	-

Purchase Price (in USD)

All Purchases	109.41 (173.51)	0.23-2800.03
<\$1	2.2% (121)	-
\$1-\$4.99	3.3% (178)	-
\$5-\$9.99	3.1% (168)	-
\$10-\$19.99	8.7% (472)	-
\$20-\$49.99	24.7% (1344)	-
\$50-\$99.99	28.2% (1532)	-
\$100-\$199.99	16.3% (884)	-
\$200-\$499.99	10.8% (589)	-
\$500-\$999.99	1.9% (201)	-
>\$1000	0.8% (44)	-

Locations Shipped From

Australia	8.74% (475)	-
Belgium	0.83% (45)	-
Belize	0.02% (1)	-
Bulgaria	0.64% (35)	-
Canada	0.61% (33)	-
China	0.02% (1)	-
Colombia	0.02% (1)	-
Czech Republic	0.09% (5)	-
Denmark	0.81% (44)	-
Europe/EU	7.19% (391)	-
France	0.74% (40)	-
Germany	25.10% (1364)	-
Hungary	0.06% (3)	-
India	0.18% (10)	-
Italy	0.99% (54)	-
Mexico	0.02% (1)	-
Netherlands	9.22% (501)	-
Norway	0.29% (16)	-
Poland	0.11% (6)	-
South Africa	0.2% (11)	-
Spain	2.37% (129)	-
Switzerland	0.39% (21)	-
UK	13.78% (749)	-
United States	19.34% (1051)	-
Unknown or N/A	8.23% (447)	-

<i>Locations Shipped To</i>		
Australia	8.19% (445)	-
Europe	15.73% (855)	-
Europe and US	0.07% (4)	-
Europe except Italy	0.18% (10)	-
Europe except UK	0.48% (26)	-
Germany	1.23% (67)	-
Switzerland	0.13% (7)	-
UK	4.42% (240)	-
United States	17.32% (941)	-
US and Canada	0.04% (2)	-
Worldwide	36.53% (1985)	-
Worldwide with exceptions	7.16% (389)	-
Unknown or N/A	8.60% (463)	-

References

- Afilipoaie, A. and Shortis, P. (2018). Crypto-Market enforcement-new strategy and tactics. *Global Drug Policy Observatory*, 54, 87-98.
- Akerlof, G. A. (1970). The market for lemons: Qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84, 488-500.
- Aldridge, J. and D. Décary-Héту (2014). “Not an ‘Ebay for Drugs’: The Cryptomarket ‘Silk Road’ as a Paradigm Shifting Criminal Innovation.” *Available at SSRN*.
- Aldridge, J. and Décary-Héту, D. (2016a) Cryptomarkets and the future of illicit drug markets. *European Monitoring Centre for Drugs and Drug Addiction* (Ed.), *The Internet and Drug Markets* (EMCDDA Insights 21), Publications Office of the European Union, Luxembourg, pp. 23-30.
- Aldridge, J. and Décary-Héту, D. (2016b). Hidden wholesale: The drug diffusing capacity of online drug cryptomarkets. *International Journal of Drug Policy*, 35, 7-15.
- Aldridge, J., Stevens, A., and Barratt, M. (2017). Will growth in cryptomarket drug buying increase the harms of illicit drugs? *Addiction*, 113(5), 789-796.
- Alm, J. and Mack, K. (2017). Degree-correlation, Robustness and Vulnerability in Finite Scale-free Networks. *Asian Research Journal of Mathematics*, 2(5), 1-6.
- Andresen, M.A., Curman, A.S., & Linning, S.J. (2017). The Trajectories of Crime at Places: Understanding the Patterns of Disaggregated Crime Types. *Journal of Quantitative Criminology*, 33, 427-449.
- Archak, N., Ghose, A., Ipeirotis, P. G. (2011). Deriving the Pricing Power of Product Features by Mining Consumer Reviews. *Management Science*, 57 (8), 1485–1509.
- Armona, L. (2017). Measuring the Impact of Formal and Informal Communication on Electronic Commerce Demand. *Stanford University mimeo*.
- Bancroft, A., and Reid, P. (2016). Concepts of illicit drug quality among darknet market users: Purity, embodied experience, craft and chemical knowledge. *International Journal of Drug Policy*, 35, 42-49.
- Barabasi, A. and Albert, R. (1999). Emergence of Scaling in Random Networks. *Science*, 286(5439), 509-512.
- Barratt, M. J. (2012). “Silk Road: eBay for Drugs.” *Addiction* 107(3): 683–683.
- Barratt, M. and Aldridge, J. (2016). Everything you always wanted to know about drug cryptomarkets* (*but were afraid to ask). *International Journal of Drug Policy*, 35, 1-6.

- Barratt, M. and Aldridge, J. (2020). No magic pocket: Buying and selling on drug cryptomarkets in response to the COVID-19 pandemic and social restrictions. *International Journal of Drug Policy*, 83, 1028-1043.
- Barratt, M., Ferris, J., and Winstock, A. (2016). Safer scoring? Cryptomarkets, social supply and drug market violence. *International Journal of Drug Policy*, 35, 24-31.
- Barratt, M., Lenton, S., Maddox, A. and Allen, M. (2016). “What if you live on top of a bakery and you like cakes?” – Drug Use and Harm Trajectories Before, During and After the Emergence of Silk Road. *International Journal of Drug Policy*, 35, 50-57.
- Barrera, V., Malm, A., Décary-Héту, D., and Munksgaard, R. (2019). Size and scope of the tobacco trade on the darkweb. *Global Crime*, 20(1), 26-44.
- Barreda, A., and Bilgihan, A. (2013). An analysis of user generated content for hotel experiences. *Journal of Hospitality and Tourism Technology*, 4(3), 263–280.
- Batiskas, M. and Kretschmer, T. (2018). Entrepreneurs on the Darknet: Reaction to Negative Feedback.
- Berk, R. (2008). How can you tell if the simulations in computational criminology are any good? *Journal of Experimental Criminology*, 4, 289–308.
- Berezina, K., Bilgihan, A., Cobanoglu, C., and Okumus, F. (2015). Understanding Satisfied and Dissatisfied Hotel Customers: Text Mining of Online Hotel Reviews, *Journal of Hospitality Marketing & Management*, 25(1), 1-24.
- Bharwani, S., & Jauhari, V. (2013). An exploratory study of competencies required to co-create memorable customer experiences in the hospitality industry. *International Journal of Contemporary Hospitality Management*, 25(6), 823–843.
- Bhaskar, A., Ramesh, D., Vichare, G., Koganti, T., Gurubaran, S. (2017). Quantitative assessment of drivers of recent global temperature variability: an information theoretic approach. *Climate Dynamics*, 49, 3877-3886.
- Bichler, G., Malm, A., and Cooper, T. (2017). Drug supply networks: a systematic review of the organizational structure of illicit drug trade. *Crime Science*, 6(2), 63-73.
- Birks, D., and Davies, T. (2017). Street network structure and crime reduction: Agent-based investigation of the encounter and enclosure hypotheses. *Criminology* 55, 900–937.
- Birks, D., Townsley, M., and Stewart, A. (2012). Generative explanations of crime: Using simulation to test criminological theory. *Criminology*, 50, 221–254.
- Bossler, A.M. and Holt, T.J. (2012), Patrol officers' perceived role in responding to cybercrime. *Policing: An International Journal*, 35(1), 165-181.
- Bossler, A. and Holt, T. (2016). On the need for policing cybercrime research. *ACJS Today*, 41(1), 14.
- Bradbury, D. (2014). Unveiling the Dark Web. *Network Security*, 4, 14-17.

- Branwen, G., Christin, N., Décary-Héту, D., Andersen, R. M., StExo, Presidente, E., Anonymous, Lau, D., Sohlz, Kratunov, D., Cakic, V., Buskirk, V., and Whom (2015). *Dark Net Market archives*, 2011-2015.
- Brenner, S. (2010). *Cybercrime: Criminal Threats from Cyberspace*. Santa Barbara, CA, Praeger.
- Bright, D. A., and Delaney, J. J. (2013). Evolution of a drug trafficking network: Mapping changes in network structure and function across time. *Global Crime*, 14, 238-260.
- Bright D., Greenhill C., Britz, T., Ritter, A., and Morselli, C. (2014) Criminal network vulnerabilities and adaptations. *Global Crime*, 18(4), 424–441.
- Bright, D. A., Greenhill, C., Levenkova, N. (2014). Dismantling criminal networks: Can node attributes play a role. In Morselli, C. (Ed.), *Crime and Networks* (pp. 148-162). New York, NY: Routledge.
- Bright, D. A., Koskinen, J., & Malm, A. (2018). Illicit network dynamics: The formation and evolution of a drug trafficking network. *Journal of Quantitative Criminology*, 35(2), 237-258.
- Broséus, J., Rhumorbarbe, D., Mireault, C., Ouellette, V., Crispino, F., Décary-Héту, D. (2017). Studying illicit drug trafficking on Darknet markets: structure and organisation from a Canadian perspective. *Forensic science International*, 264, 7–14.
- Burt, R.S., (2000). ‘The Network Structure of Social Capital’, in R. I. Sutton & B. M. Staw, (eds). *Research in Organizational Behavior*, 22 (May), pp.345–423.
- Cabral, L. and Hortac, su A. (2010). The Dynamics of Seller Reputation: Evidence from e-Bay. *Journal of Industrial Economics*, 58, 54–78.
- Campana, P. and Varese, F. (2013). Cooperation in Criminal Organizations: Kinship and Violence as Credible Commitments. *Rationality and Society*, 25(3), 263-289.
- Cantalops, A. S., and Salvi, F. (2014). New consumer behavior: a review of research on eWOM and hotels. *Int. J. Hosp. Manage.* 36, 41-51.
- Carley, K.M. (1995). Computational and Mathematical Organization Theory: Perspective and Directions. *Computational and Mathematical Organization Theory*, 1(1), 39–56.
- Casson, M. (2001). *Information and Organization: A New Perspective on the Theory of the Firm*. New York: Oxford University Press.
- Catino, M. (2014). How Do Mafias Organize?: Conflict and Violence in Three Mafia Organizations. *European Journal of Sociology*, 55(2), 177-220.
- Chalmers. D. and Bradford, J. (2013) Methamphetamine users’ perceptions of exchanging drugs for money: Does trust matter?. *Journal of Drug Issues*, 43(3), 256-269.
- Christin, N. (2013). *Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace*. Proceedings of the 22nd international conference on World Wide Web, International World Wide Web Conferences Steering Committee.
- Chow, S. F., and Zhang, L. L. (2008). Measuring consumer satisfaction and dissatisfaction intensities to identify satisfiers and dissatisfiers. *Journal of Consumer Satisfaction, Dissatisfaction and Complaining Behavior*, 21, 66–77.

- Coase, R. (1937). The Nature of the Firm. *Economica*, 4(16), 386-405.
- Cohen, L. and Felson, M. (1979). Social Change and Crime Rate Trends: A Routine Activity Approach. *American Sociological Review*, 44(4), 588-608.
- Cook, K. and Emerson, M. (1978). Power, Equity and Commitment in Exchange Networks. *American Sociological Review*, 43, 721-739.
- Cox, J. (2016). Staying in the shadows: The use of bitcoin and encryption in cryptomarkets EMCDDA (Ed.), *Internet and drug markets*, EMCDDA insights, Publications Office of the European Union, Luxembourg, 41-47.
- Cressey, D. (1967). Methodological Problems in the Study of Organized Crime as a Social Problem. *AAPSS*, 374(1), 101-112.
- Dasgupta, N., Freifeld, C., Brownstein, J., Menone, C., Surratt, H., Poppish, L., Green, J., Lavonas, E., and Dart, R. (2013). Crowdsourcing Black Market Prices for Prescription Opioids. *Journal of Medical Internet Research*, 15(8).
- Decary-Hetu, D. (2016). Policing cybercrime and cyberterror. *Global Crime*, 17(1), 123-125.
- Decary-Hetu, D. and Dupont, B. (2012). The social network of hackers. *Global Crime*, 13(3), 1-16.
- Decary-Hetu, D., and Giommoni, L. (2017). Do police crackdowns disrupt drug cryptomarkets? A longitudinal analysis of the effects of Operation Onymous. *Crime, Law, and Social Change*, 67(1), 55-75.
- Décary-Héту, D., Mousseau, V., and Rguioui, I. (2018). The Shift to Online Tobacco Trafficking. *International Journal of Cyber Criminology*, 12(1), 47-67.
- Décary-Héту, D., Mousseau, V., and Vidal, S. (2018). Six Years Later: Analyzing Online Black Markets Involved in Herbal Cannabis Drug Dealing in the United States. *Contemporary Drug Problems*, 45(4), 366-381.
- Décary-Héту, D. and Quessy-Dore, O. (2017). Are Repeat Buyers in Cryptomarkets Loyal Customers? Repeat Business Between Dyads of Cryptomarket Vendors and Users. *American Behavioural Scientist*, 61(11), 1341-1357.
- Decker, R. and Trusov, M. (2010). Estimating aggregate consumer preferences from online product reviews. *International Journal of Research in Marketing*, 27, 293-307.
- Demant, J., Munksgaard, R., Houborg, E. (2018). Personal use, social supply or redistribution? Cryptomarket demand on silk road 2 and Agora. *Trends in Organized Crime*, 21(1), 42-61.
- Diekmann A., Jann, B., Przepiorka, W., and Wherli, S. (2014). Reputation formation and the evolution of cooperation in anonymous online markets. *Am Soc Rev*, 79, 65-85.
- Dingledine, R., et al. (2004). Tor: The Second-Generation Onion Router, DTIC Document.

- Dittus, WPJ. (2017). The biogeography and ecology of Sri Lankan mammals point to conservation priorities. *Ceylon Journal of Science*, 46, 33–64.
- Dolliver, D., Ericson, S., and Love, K. (2018) A Geographic Analysis of Drug Trafficking Patterns on the TOR Network. *Geographical Review*, 108(1), 45-68.
- Dolliver, D. and Kuhns, J. (2016). The Presence of New Psychoactive Substances in a Tor Network Marketplace Environment. *J Psychoactive Drugs*, 48(5), 321-329.
- Duijn, P., Kashirin, V., & Sloot, P. (2014). The relative ineffectiveness of criminal network disruption. *Scientific Reports*, 4, 4238.
- Dumouchel, P. (2005). Trust as an Action. *European Journal of Sociology*, 46(3), 417-428.
- Dupont, B., Côté, A., and Savine, C., and Décary-Héту, D. (2016). The Ecology of Trust Among Hackers. *Global Crime*, 17(2), 129-151.
- Duxbury, S. and Haynie, D. (2017). The Network Structure of Opioid Distribution on a Darknet Cryptomarket. *Journal of Quantitative Criminology*, 34(4), 921-941.
- Duxbury, S. and Haynie, D. (2018). Building them up, breaking them down: Topology, vendor selection patterns, and a digital drug market's robustness to disruption. *Social Networks*, 52, 238-250.
- Duxbury, S. and Haynie, D. (2019). Criminal network security: An agent-based approach to evaluating network resilience. *Criminology*, 57(2), 314-342.
- Duxbury, S. and Haynie, D. (2020). The responsiveness of criminal networks to intentional attacks: Disrupting darknet drug trade. *Plos one*, 15(9), 1-20.
- Eck, J., Clarke, R., & Guerette, R. (2007). Risky facilities: Crime concentrations in homogeneous sets of establishments and facilities. *Crime Prevention Studies*, 21, 255–264.
- Ekinci, Y., Dawes, P. L., & Massey, G. R. (2008). An extended model of the antecedents and consequences of consumer satisfaction for hospitality services. *European Journal of Marketing*, 42(1/2), 35–68.
- Eliashberg, J., Hui, S., and Zhang, J. (2007). From Story Line to Box Office: A New Approach for Green-Lighting Movie Scripts. *Management Science*, 53(6), 881- 893.
- Farley, M., Franzblau, K., and Kennedy, M. (2013). Online Prostitution and Trafficking. *Albany Law Review*, 77(3), 1039-1094.
- Foley, S., Karlsen, J., and Putnins, T. (2018). Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed Through Cryptocurrencies?.. *Review of Financial Studies*, 32(5), 1798-1853.
- Gambetta, D. (1988). *Trust: Making and Breaking Cooperative Relations*. Blackwell.

- Gambetta, D. (1993). *The Sicilian Mafia; The Business of Private Protection*. Cambridge, MA: Harvard University Press.
- Gambetta, D. (2000). Can we trust trust? In: D. Gambetta, ed. *Trust: Making and Breaking Cooperative Relations*. Oxford: Department of Sociology, University of Oxford, 213-237.
- Gambetta, D. (2009). *Codes of the Underworld: How Criminals Communicate*. Princeton, NJ: Princeton University Press.
- Gambetta, D. and Bacharach, M. (2001). Trust in Signs. In K. Cook (ed.) *Trust and Society*, New York: Russell Sage Foundation, 148–184.
- Gaup Moe, Marie Elizabeth. (2008). Quantification of Anonymity for Mobile Ad Hoc Networks. *Electronic Notes in Theoretical Computer Science*, 1-12.
- Genolini, C. & Falissard, B. (2010). KmL: k-means for longitudinal data. *Comput Stat*, 25(2), 317–328.
- Ghose, A., Ipeirotis, P. G. (2011). Estimating the Helpfulness and Economic Impact of Product Reviews: Mining Text and Reviewer Characteristics. *IEEE Transactions on Knowledge & Data Engineering*, 23(10), 1498–1512.
- Gilbert, N., & Troitzsch, K. G. (2005). *Simulation for the social scientist*. Open University Press.
- Godes, D. and Silva, J. C. (2012). Sequential and temporal dynamics of online opinion. *Marketing Science*, 31(3), 448–473.
- Goeman, J. J. (2010). L1 penalized estimation in the cox proportional hazards model. *Biometrical Journal*, 52(1), 70–84.
- Goodman, M. (1997). Why the Police Don't Care about Computer Crime. *Harvard Journal of Law & Technology*, 10(3), 465-494.
- Grabosky, P. (2001). Virtual Criminality: Old Wine in New Bottles? *Social & Legal Studies*, 10(2), 243-249.
- Grabosky, P. (2007). The Internet, Technology, and Organized Crime. *Asian Criminology*, 2, 145-161.
- Grabosky, P. N., Smith, R. (2001). Telecommunication fraud in the digital age: The convergence of technologies. In Wall, D. (Ed.), *Crime and the Internet* (pp. 29-43). New York, NY: Routledge.
- Greif, A. (1989). Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders. *The Journal of Economic History*, 49(4), 857–882.
- Groff, E. R., Johnson, S. D., & Thornton, A. (2018). State of the art in agent-based modeling of urban crime: An overview. *Journal of Quantitative Criminology*, 35, 155-193.

- Hamill H (2011) *The Hoods: Crime and Punishment in West Belfast*. Princeton, NJ: Princeton University Press.
- Hardy, R. and Norgaard, J. (2016). Reputation in the Internet black market: an empirical and theoretical analysis of the Deep Web. *Journal of Institutional Economics*, 12(3), 515-539.
- Herley, C. and Florencio, D. (2010). “Nobody Sells Gold for the Price of Silver: Dishonesty, Uncertainty and the Underground Economy”. In *Economics of Information Security and Privacy*, Edited by: Moor, Tyler, Pym, David J. and Ioannidis, Christos. 35–53. New York: Springer.
- Hillman, H. and Aven, B. (2011). Fragmented Networks and Entrepreneurship in Late Imperial Russia. *American Journal of Sociology*, 117(2), 484-538.
- Hinduja, S. and Patching, J. (2012). Cyberbullying: Neither an epidemic nor a rarity. *European Journal of Developmental Psychology*, 9(5), 539-543.
- Hollenbeck, Brett (2017). The Economic Advantages of Chain Affiliation. *RAND Journal of Economics*, 48 (4), 1103–35.
- Holt, T., Bossler, A., and Malinski, R. (2016). Identifying Predictors of Unwanted Online Sexual Conversations Among Youth Using a Low Self-Control and Routine Activity Framework. *Journal of Contemporary Criminal Justice*, 32(2), 108-128.
- Holt, T. and Lampke, E. (2010). Exploring stolen data markets online: products and market forces. *Criminal Justice Studies*, 23(1), 33-50.
- Holt, T., Strumsky, D., Smirnova, O., and Kilger, M. (2012). Examining the social networks of malware writers and hackers. *International Journal of Cyber Criminology*, 6(1), 891.
- Holt, T. and Turner, M. (2012). Examining risks and protective factors of on-line identity theft. *Deviant Behavior*, 33(4), 308-323.
- Hutchings, A. & Holt, T. J. (2015). A crime script analysis of the online stolen data market. *British Journal of Criminology*, 55(3), 596-614.
- Ianelli, N. and Hackworth, A. (2005). Botnets as a Vehicle for Online Crime. *Carnegie Mellon University*, 1-30.
- Jacques, S., Allen, A., and Wright, R. (2014). Drug dealers’ rational choices on which customers to rip-off. *International Journal of Drug Policy*, 25(2), 251-256.
- Janze, C. (2017). Are Cryptocurrencies Criminals Best Friends? Examining the Co-Evolution of Bitcoin and Darknet Markets. *Twenty-third Americas Conference on Information Systems*.
- Janetos, N. and Tilly, J. (2017). Reputation Dynamics in a Market for Illicit Drugs. *Unpublished Paper*.
- Jappelli, T. and Pagano, M. (2002). Information sharing, lending and defaults: Cross-country evidence. *Journal of Banking and Finance*, 26(10), 2017-2045.

Kenney, M. (2007). The Architecture of Drug Trafficking: Network Forms of Organisation in the Colombian Cocaine Trade. *Global Crime*, 8(3), 233–259.

Kennedy, D. (2008). *Deterrence and crime prevention: Reconsidering the prospect of sanction*. London, England: Routledge.

Klerks, P. (2001). The network paradigm applied to criminal organisations: Theoretical nitpicking or a relevant doctrine for investigators?. *Recent developments in the Netherlands. Connections*, 24(3), 53–65.

Kowalski, M., Hooker, C., and Barratt, M. (2019). Should we smoke it for you as well? An ethnographic analysis of a drug cryptomarket environment. *International Journal of Drug Policy*, 73, 245-254.

Krebs, V. (2001). Mapping networks of terrorist cells. *Connections*, 24(3), 43–52.

Keegan, B., Ahmed, M., Williams, D., Srivastava, J., and Contractor, N. 2010. Dark Gold: Statistical Properties of Clandestine Networks in Massively Multiplayer Online Games. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*. 201–208.

Kennedy, D. (2008). *Deterrence and crime prevention: Reconsidering the prospect of sanction*. London, England: Routledge.

Khodyakov, D., (2007). ‘Trust as a Process’. *Sociology*, 41(1), 115–132.

Kollock, P. (1999). The Production of Trust in Online Markets. *Advances in Group Processes*, 99-123.

Ladegaard, I. (2017). We Know Where You Are, What You Are Doing and We Will Catch You: Testing Deterrence Theory in Digital Drug Markets. *The British Journal of Criminology*, 58(2), 414-433.

Lamy, F., Daniulaityte, R., Barratt, M., Lokala, U, Sheth, A., and Carlson R. (2020). Listed for sale: Analyzing data on fentanyl, fentanyl analogs and other novel synthetic opioids on one cryptomarket. *Drug and Alcohol Dependence*, 213, 1-8.

Lacson, W. & Jones, B. (2016). The 21st Century DarkNet Market: Lessons from the Fall of Silk Road. *International Journal of Cyber Criminology*, 10 (1), 40-61.

Lee, T. and Bradlow, E. (2011). Automated Marketing Research Using Online Customer Reviews. *Journal of Marketing Research*, 48(5), 881-894.

Leeson, P. T. (2005). Endogenizing Fractionalization. *Journal of Institutional Economics*, 1(1): 75–98.

Lewman, A. (2016). Tor and links with cryptomarkets EMCDDA (Ed.), *Internet and drug markets*, EMCDDA insights, Publications Office of the European Union, Luxembourg, 33-40.

Litvin, S. W., Goldsmith, R. E., and Pan, B. (2008). Electronic word-of-mouth in hospitality and tourism management. *Tour. Manage*, 29, 458–468.

- Lo, T.W., (2010). 'Beyond Social Capital: Triad Organized Crime in Hong Kong and China'. *British Journal of Criminology*, 50 (5), 851–872.
- Luca, M. (2011): "Reviews, Reputation, and Revenue: The Case of Yelp.com". *mimeo*.
- Lusthaus, J. (2012). Trust in the world of cybercrime. *Global crime* 13 (2), 71-94
- Lusthaus, J. (2018). *Industry of Anonymity*. Cambridge: Harvard University Press.
- Maddox, A., Barratt, M., Allen, M., and Lenton, S. (2016). Constructive activism in the dark web: Cryptomarkets and illicit drugs in the digital 'demimonde' Information. *Communication and Society*, 19, 111-126.
- Malm, A. and G. Bichler (2011). Networks of Collaborating Criminals: Assessing the Structural Vulnerability of Drug Markets. *Journal of Research in Crime and Delinquency*, 48(2), 271-297.
- Mathewson, N. and Dingedine, R. (2004). Practical Traffic Analysis: Extending and Resisting Statistical Disclosure. *Privacy Enhancing Technologies Workshop*.
- Martin, J. (2013). "Lost on the Silk Road: Online Drug Distribution and the "Cryptomarket". *Criminology and Criminal Justice*, 14(3): 351-367.
- Martin, J. (2014a). *Drugs on the Dark Net*. New York, NY: Palgrave Macmillan.
- Martin, J. (2014b). Lost on the Silk Road: Online Drug Distribution and the "Cryptomarket", *Criminology & Criminal Justice*, 14(3), 351-367.
- Martin, J., Cunliffe, J., Decary-Hetu, D., and Aldridge, J. (2018). Effect of restricting the legal supply of prescription opioids on buying through online illicit marketplaces: interrupted time series analysis. *BMJ*, 361, 1-27.
- May, T. and Hough, M. (2004). Drug markets and distribution systems. *Addiction Research & Theory*, 12(6), 549-563.
- McGloin, J. M. (2005). Policy intervention and the considerations of a network analysis of street gangs. *Criminology & Public Policy*, 4, 607–636.
- McGloin, J. M. and Kirk, D. (2011). An overview of social network analysis. *Journal of Criminal Justice Education*, 2(2), 169-181.
- McGloin, J. M., and Rowan, Z. (2015). A threshold model of collective crime. *Criminology*, 53, 484–512.
- Mikroyannidis, A. and Theodoulidis, B. (2006). Heraclitus ii: A Framework for Ontology Management and Evolution in Proceedings of the 2006 *IEEE/WIC/ACM International Conference on Web Intelligence*, Nishida, Toyooki (Chair). Hong Kong, China: IEEE Computer Society, 514–521.

- Milgrom, P., North, D., and Weingast, B. (1990). The Role of Institutions in the Revival fo Trade: The Law Merchant, Private Judges, and the Champagne Friars. *Economics and Politics*, 2(1), 1-23.
- Mohamed, R. and Fritsvold, E. (2011). Is the College Campus a Safe Haven for Drug Dealers?: Dorm Room Dealers: Drugs and the Privileges of Race and Class. *Symbolic Interaction*, 34(2), 309-311.
- Morselli, C. (2009). *Inside criminal networks*. New York: Springer.
- Morselli, C., Decary-Hetu, D., Paquet-Clouston, M., and Aldridge, J. (2017). Conflict Management in Illicit Drug Cryptomarkets. *International Criminal Justice Review*, 27(4), 237-254.
- Morselli, C., Giguere, C., and Petit, K. (2007). The efficiency/security trade-off in criminal networks. *Social Networks*, 29(1), 143–153.
- Munksgaard, R. and Demant, J. (2016). Mixing politics and crime – The prevalence and decline of political discourse on the cryptomarket. *International Journal of Drug Policy*, 35, 77-83.
- Nagin, D. & Land, K. (1993). Age, criminal careers and population heterogeneity: specification and estimation of a nonparametric, mixed Poisson model. *Criminology*, 31(3), 327–362.
- Natarajan, M. (2006). Understanding the structure of a large heroin distribution network: A quantitative analysis of qualitative data. *Journal of Quantitative Criminology*, 22(2), 171-192.
- Newman, M.E.J. (2002). Assortative mixing in networks. *Phys. Rev. Lett.*, 89, 2087-3001.
- Newman, M.E.J. (2003). The structure and function of complex networks. *SIAM Rev.*, 45, 167–256.
- Newman, M. (2006). Modularity and community structure in networks. *Proc. Natl. Acad. Sci. U.S.A.* 103(23), 8577-8582.
- Newman, M. and Girvan, M. (2004). Finding and evaluating community structure in networks. *Phys. Rev. E*, 69, 26-113.
- Norbutas, L. (2018). Offline constraints in online drug marketplaces: An exploratory analysis of a cryptomarket trade network. *International Journal of Drug Policy*, 56, 92-100.
- Norbutas, L., Ruiters, S., and Corten, R. (2020). Believe it when you see it: Dyadic embeddedness and reputation effects on trust in cryptomarkets for illegal drugs. *Social Networks*, 63, 150-161.
- Norgaard, J., Walbert, H., and Hardy, R. (2018). Shadow markets and hierarchies: comparing and modeling networks in the Dark Net. *Journal of Institutional Economics*, 1-23.

- Ormsby, E. (2013). *The Drug's in the Mail*. *The Age*. Melbourne, Fairfax.
- Papachristos, A. V. (2009). Murder by structure: Dominance relations and the social structure of gang homicide. *American Journal of Sociology*, *115*, 74–128.
- Papachristos, A. V. (2011). The coming of a networked criminology. *Advances in Criminological Theory*, *17*, 101–140.
- Papachristos, A. V. (2014). The network structure of crime. *Sociology Compass*, *8*, 347–357.
- Paquet-Clouston, M., Décary-Héту, D., and Bilodeau, O. (2018). Cybercrime is where responsibility is a case study of an online behavior system in crime. *Global Crime*, *19* (1), 1-21.
- Park, D.-H., Lee, J., and Han, I. (2007). The effect of on-line consumer reviews on consumer purchasing intention: The moderating role of involvement. *International Journal of Electronic Commerce*, *11*(4), 125-148.
- Pease, Ken. 1991. The Kirkholt Project: Preventing Burglary on a British Public Housing Estate. *Security Journal* *2*, 73-77.
- Pekar, V. and Ou, S. (2008). Discovery of subjective evaluations of product features in hotel reviews. *J Vacat Mark*, *14*(2):145–155.
- Pizam, A., and Ellis, T. (1999). Customer satisfaction and its measurement in hospitality enterprises. *International Journal of Contemporary Hospitality Management*, *11*(7), 326–339.
- Pons, P. and Latapy, M. (2005). Computing communities in large networks using random walks. *International Symposium on Computer and Information Sciences*, Springer, 284-293.
- Prentice, C. (2013). Service quality perceptions and customer loyalty in casinos. *International Journal of Contemporary Hospitality Management*, *25*(1), 49–64.
- Przepiorka, W., and Aksoy, O. (2017). *Social Order in Online Markets and the 'Collapse' of Institutions*. Unpublished manuscript, The Netherlands: Department of Sociology/ICS, Utrecht University.
- Przepiorka, W., Norbutas, L., and Corten, R. (2017). Order without Law: Reputation Promotes Cooperation in a Cryptomarket for Illegal Drugs. *European Sociological Review*, *33*(6), 752-764.
- Qin, J., Xu, J., Hu, D., Sageman, M., and Chen, H. (2005). Intelligence and security informatics. In: *Analyzing Terrorist Networks: A Case Study of the Global Salafi Jihad Network*. Springer, Berlin, Heidelberg, pp. 287–304.
- Raab, J. and Milward, B. (2003). Dark networks as problems. *J Public Adm Res Theory*, *13*(4), 413-439.

- Resnick P., Zeckhauser R. (2002). Trust among strangers in internet transactions: empirical analysis of eBay's reputation system. In Baye M. R. (Ed.), *The Economics of the Internet and E-Commerce*. Amsterdam: Elsevier, pp. 127–157.
- Reuter P. (1985). *Disorganized Crime*. Chapter 6. “Violence and Market Organization”: 132-150.
- Savona, E. and Mignone, M. (2004). The Fox and the Hunters: How IC Technologies Change the Crime Race. In *Crime and Technology: New Frontiers for Regulations, Law Enforcement and Research* (7-28).
- Salt, J. and Stein, J. (2002). Migration as a Business: The Case of Trafficking. *International Migration*, 35(4), 467-494.
- Schwartz, D. and Rouselle, D. (2008). Targeting criminal networks: Using social network analysis to develop enforcement and intelligence priorities. *IALEIA J*, 18(1), 18–44.
- Seshadri, T. and Tellis, G. (2012). Does chatter really matter? Dynamics of user-generated content and stock performance Mark. *Science*, 31(2), 198–215.
- Shapiro C. (1983). Premiums for high quality products as return to reputation. *Quarterly Journal of Economics*, 98, 659–680.
- Shaw, C. R., & McKay, H. D. (1942). *Juvenile delinquency and urban areas*. Chicago, IL: University of Chicago Press.
- Shearing, C. and Wood, J. (2003). Nodal Governance, Democracy, and the New ‘Denizens’. *Journal of Law and Society*, 30(3), 400-419.
- Sherman, L. (2007). The power few: experimental criminology and the reduction of harm. *Journal of Experimental Criminology*, 3, 299-321.
- Shortis, P., Aldridge, J., and Barratt, M. J. (2020). Drug cryptomarket futures: Structure, function and evolution in response to law enforcement actions. In *Research Handbook on International Drug Policy* (pp. 355-379). Edward Elgar Publishing Ltd.
- Simon, H. (1955). On a class of skew distribution functions. *Biometrika*, 42, 425–440.
- Soska, K. and Christin, N. (2015). Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. Paper presented at the *24th USENIX Security Symposium*, Washington, D.C.
- Sparrow, MK. (1991). Network vulnerabilities and strategic intelligence in law enforcement. *J of Intell and Counterintell*, 5(3), 255–274.
- Stephen, A. and Toubia, O. (2009). Explaining the power-law degree distribution in a social
- Szumilas, M. (2010). Explaining odds ratios. *Journal of the Canadian Academy of Child and Adolescent Psychiatry*, 19(3), 227-229.
- Taboada, M., Brooke, J., Tofiloski, M., Voll, K. and Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37, 267-307.

- Thomaz, F., Salge, C., Karahanna, E., and Hulland, J. (2020). Learning from the Dark Web: leveraging conversational agents in the era of hyper-privacy to enhance marketing. *Journal of the Academy of Marketing Science*, 48(1), 43-63.
- Torres, E., & Kline, S. (2013). From customer satisfaction to customer delight: Creating a new standard of service for the hotel industry. *International Journal of Contemporary Hospitality Management*, 25(5), 642–659.
- Tsvetovat, Max & Carley, Kathleen. (2003). Bouncing Back: Recovery mechanisms of covert networks. *NAACSOS Conference 2003*.
- Tsuchiya, Y. and Hiramoto, N. (2021). Dark Web in the Dark: Investigating when Transactions Take Place on Cryptomarkets. *Forensic Science International: Digital Investigation*, 1-38.
- Tzanetakis, M., Kamphausen, G., Werse, B., and von Laufenberg, R. (2015). The transparency paradox. Building trust, resolving disputes and optimising logistics on conventional and online drugs markets. *International Journal of Drug Policy*, 35, 58-68.
- Ur-Rahman, N. and Harding, J. (2011). Textual Data Mining for Industrial Knowledge Management and Text Classification: A Business Oriented Approach. *Expert Systems with Applications*, 39 (5), 1–11.
- Van Hout, M, and Bingham, T. (2013a). ‘Surfing the Silk Road’: A study of users’ experiences. *International Journal of Drug Policy*, 24, 524-529.
- Van Hout, M, and Bingham, T. (2013b). 'Silk Road' The virtual drug marketplace: A single case study of user experiences. *International Journal of Drug Policy*, 24(5), 155-165.
- Van Hout, M, and Hearne, E. (2016). New psychoactive substances (NPS) on cryptomarket fora: An exploratory study of characteristics of forum activity between NPS buyers and vendors. *International Journal of Drug Policy*, (40), 102-110.
- Van Buskirk, J., Bruno, R., Dobbins, T., Breen, C., Burns, L., Naicker, S., and Roxburgh, A. (2017). The recovery of online drug markets following law enforcement and other disruptions. *Drug and Alcohol Dependence*, 173, 159-162.
- van Der Heide, B., Johnson, B. K., & Vang, M. H. (2013). The effects of product photographs and reputation systems on consumer behavior and product cost on eBay. *Computers in Human Behavior*, 29(3), 570–576.
- van Wegberg, R. and Verburgh, T. (2018). Lost in the dream? Measuring the effects of operation bayonet on vendors migrating to dream market. *Proceedings of the evolution of the darknet workshop*, 1-5.
- Varese, F. (2010). Introduction to F.V. (ed.), *Organized Crime* (series: Critical Concepts in Criminology), 4 vols. *Routledge*, 1-35.
- Varese F. (2010). General Introduction: What is Organized Crime?. in Federico Varese (ed.) *Organized Crime: Critical Concepts in Criminology*, New York: Routledge, 1-35.

- Vayrynen, R. (2003). Illegal Immigration, Human Trafficking, and Organized Crime. *WIDER*, 1-25.
- Venkatesh, V. and Goyal, S. (2010). Expectation disconfirmation and technology adoption: polynomial modeling and response surface analysis. *MIS Quarterly*, 34(2), 281-303.
- Volery, T., Mueller, S., and von Siemens, B. (2013). Entrepreneur ambidexterity: A study of entrepreneur behaviours and competencies in growth-oriented small and medium-sized enterprises. *International Small Business Journal*, 33(2), 109-129.
- Von Lampe, K. (2016). *Organized Crime*. New York, NY: Sage Publications.
- Von Lampe, K. and Johansen, P. (2004). Organized Crime and Trust: On the conceptualization and empirical relevance of trust in the context of criminal networks. *Global Crime*, 6(2), 159-184.
- Wagner, R. (2017), 'Does feedback to business-plans impact new ventures? fundraising in a randomized field experiment'. Working Paper.
- Wall, D. S. (2001). Maintaining order and law on the internet. In D. S. Wall (Ed.), *Crime and the internet* (pp. 167–183). London: Routledge.
- Wall, D. (2007). Policing Cybercrimes: Situating the Public Police in Networks of Security within Cyberspace. *Police Practice and Research*, 8(2), 183–205.
- Wall, D. (2015). The Internet as a conduit for criminal activity. In Pattavina, A. (Ed.), *Information technology and the criminal justice system* (77-98).
- Warr, M. (2002). *Companions in crime: The social aspects of criminal conduct*. Cambridge, England: Cambridge University Press.
- Wasserman, S. and Faust, K. (1994). *Social network analysis: methods and applications*. Cambridge University Press: Cambridge.
- Weisburd, D., Braga, A. A., Groff, E. R., & Wooditch, A. (2017). Can hot spots policing reduce crime in urban areas? An agent-based simulation. *Criminology*, 55(1), 137-173.
- Westlake, B., Bouchard, M. and Frank, R. (2011). Finding the Key Players in Online Child Exploitation Networks. *Policy & Internet*, 3(2), 1-32.
- Whitmeyer, J. M. (2000), Power Over Groups Through Effective Monitoring and Sanctioning, conference presentation: *American Sociological Association*.
- Williamson, O. (1993). Calculativeness, Trust, and Economic Organization. *Journal of Law and Economics*, 36(1), 453-486.
- Wilkins, H., Merrilees, B., & Herington, C. (2007). Towards an understanding of total service quality in hotels. *International Journal of Hospitality Management*, 26(4), 840–853.
- Wood, G. (2017). The structure and vulnerability of a drug trafficking collaboration network. *Social Network*, 48, 1-9.
- Wright, R., & Decker, S. (1994). *Burglars on the Job: Streetlife and Residential Break-ins*. Boston, MA: Northeastern University Press.

- Xu, S. and Chen, T. (2003). Robust filtering for uncertain stochastic time-delay systems. *Asian Journal of Control*, 5(3), 364–373.
- Yamagishi, T. and Matsuda, M. (2003). The Role of Reputation in Open and Closed Societies: An Experimental Study of Online Trading. *Center for the Study of Cultural and Ecological Foundations of Mind, Working Paper Series 8*.
- Yang, S., Keller, F., and Zheng, L. (2013). *Social Network Analysis: Methods and Examples*. New York: Sage Publications.
- Yar, M. (2005). The novelty of cybercrime. *European Journal of Criminology*, 2, 407-427.
- Yip, M., Webber, C. & Shadbolt, N. (2013). Trust Among Cybercriminals? Carding Forums, Uncertainty and Implications for Policing. *Policing & Society*, 23(4), 516-539.
- Zajác, R. (2017). Silk Road: The market beyond the reach of the state. *The Information Society*, 33(1), 23-34.
- Zhang, J., Meeson, A., Welchman, AE., and Kourtzi, Z. (2010). Learning alters the tuning of functional magnetic resonance imaging patterns for visual forms. *J Neurosci* 30, 14127–14133.
- Zipf, G. (1949). *Human behavior and the principle of least effort: An introduction to human ecology*. Cambridge, MA: Addison-Wesley Press.

Acknowledgments

First and foremost, I would like to thank my supervisor, Paolo, and advisor, Alice. Your guidance and support were invaluable throughout this process. You are the real MVPs. A special thank you to Daniel for his help in the early going. I still owe you dinner, pal.

To James, Helen, Godfrey, and Honor, I got 'er done. Your extreme kindness cannot be repaid. Thank you for welcoming me into your home and accepting me as one of your own. I would have certainly complained more about Cambridge were it not for you. Thank you for keeping me grounded.

To Matt, Aliyah, and Dennis, thank you for your support. To Veronica, words cannot describe my thanks.

Thank you to Rob, Steve, Pat, and Noah, my mournival of morons. Good friends are hard to come by. Our dinners were often the highlight of my week. All good conspiracies require great conspirators.

To Jason, Robin, and Bethany, much is owed to your friendship and mentorship. I would not be where I am were it not for your guidance.

To my brother, Zeljko, words cannot adequately express my appreciation. Loyalty is worth its weight in gold. You have been and always will be the man fighting beside me in the trenches. We ride or die.

To Titian, we haven't always seen eye-to-eye, but you've always had my best interest. Together, we rose from the dirt. This is as much yours as it is mine.

And finally, thanks be to the Lord above. Nothing worth having comes easy. Everything of value must be earned. Mistakes were made. Failures were had. But I have learned and improved.